Published in final edited form as:

Psychol Rev. 2013 October ; 120(4): 873–902. doi:10.1037/a0034247.

Discrete-Slots Models of Visual Working-Memory Response Times

Christopher Donkin,

University of New South Wales

Robert M. Nosofsky,

Indiana University Bloomington

Jason M. Gold, and

Indiana University Bloomington

Richard M. Shiffrin

Indiana University Bloomington

Abstract

Much recent research has aimed to establish whether visual working memory (WM) is better characterized by a limited number of discrete all-or-none slots or by a continuous sharing of memory resources. To date, however, researchers have not considered the response-time (RT) predictions of discrete-slots versus shared-resources models. To complement the past research in this field, we formalize a family of mixed-state, discrete-slots models for explaining choice and RTs in tasks of visual WM change detection. In the tasks under investigation, a small set of visual items is presented, followed by a test item in 1 of the studied positions for which a change judgment must be made. According to the models, if the studied item in that position is retained in 1 of the discrete slots, then a memory-based evidence-accumulation process determines the choice and the RT; if the studied item in that position is missing, then a guessing-based accumulation process operates. Observed RT distributions are therefore theorized to arise as probabilistic mixtures of the memory-based and guessing distributions. We formalize an analogous set of continuous shared-resources models. The model classes are tested on individual subjects with both qualitative contrasts and quantitative fits to RT-distribution data. The discrete-slots models provide much better qualitative and quantitative accounts of the RT and choice data than do the shared-resources models, although there is some evidence for “slots plus resources” when memory set size is very small.

© 2013 American Psychological Association

Correspondence concerning this article should be addressed to Christopher Donkin, Psychology Department, Mathews Building, University of New South Wales, Kensington, Sydney 2052, New South Wales, Australia, or to Robert Nosofsky, Department of Psychological and Brain Sciences, Indiana University Bloomington, Bloomington, IN47405. christopher.donkin@gmail.com or nosofsky@indiana.edu.

Christopher Donkin, Psychology Department, University of New South Wales, Kensington, New South Wales, Australia; Robert M. Nosofsky, Jason M. Gold, and Richard M. Shiffrin, Department of Psychological and Brain Sciences, Indiana University Bloomington.

Christopher Donkin and Robert M. Nosofsky contributed equally to this research and are listed alphabetically.

Keywords

visual working memory; response times; mathematical modeling

Visual working memory (WM) is the short-term memory system that maintains visual representations of stimulus inputs. It serves as a foundation for numerous cognitive processes and tasks, including the ability to locate targets embedded in distractors, to comprehend and reason about visual displays, and to detect changes in visual scenes.

An ongoing theoretical debate concerns decision making in visual WM: Is it best characterized in terms of discrete “slots,” each with all-or-none properties, or in terms of resources shared in a more continuous fashion across a set of to-be-remembered items (e.g., Alvarez & Cavanaugh, 2004; Awh, Barton, & Vogel, 2007; Barton, Ester, & Awh, 2009; Bays, Catalao, & Husain, 2009; Bays, Gorgoraptis, Wee, Marshall, & Husain, 2011; Bays & Husain, 2008; Cowan, 2001; Cowan & Rouder, 2009; Luck & Vogel, 1997; Pashler, 1988; Rouder et al., 2008; van den Berg, Shin, Chou, & George, 2012; Vogel, Woodman, & Luck, 2001; Wilken & Ma, 2004; Zhang & Luck, 2008)?

According to the discrete-slots view, visual WM makes available some number of slots for storing to-be-remembered items. The slot-based memories are conceptualized as being all-or-none: When memory is probed, if the test item occupies one of the discrete slots, then the observer can judge its presence with no loss in resolution, regardless of the number of other items in the set of to-be-remembered objects. By contrast, if the object has not been stored in one of the discrete slots, then there is a complete loss of resolution, that is, no information about the presented object remains. Specific members of the family of discrete-slots models differ according to whether the number of slots is presumed to be fixed or variable across trials and/or conditions. The central focus of the present investigation, however, will be on the mixed-state, all-or-none aspect of these models. Thus, we use the term “discrete-slots” models, rather than the more common term “fixed-slots” models, to refer to members of this family.

An alternative view proposes that visual WM consists of a pool of resources that is allocated in continuous fashion through sharing of the resources. Thus, if the number of memory items is small, then the observer can maintain high-resolution representations of all of them. By contrast, if a large number of items must be maintained, then the continuous sharing of resources leads the observer to have lower-resolution representations of the individual items when memory is probed.

Hybrid models have also been proposed in which there is sharing of resources up to a fixed-number limit, but then complete loss of resolution for any item that is not stored in one of the discrete slots (e.g., Cowan & Rouder, 2009; Zhang & Luck, 2008).

These competing views of visual WM have led to an extensive debate in the field. However, the extant research has contrasted the models only on their predictions of various forms of accuracy data. Among the major types of accuracy data are receiver operating characteristic

(ROC) curves, which plot hit rates against false alarm rates as a function of experimental conditions, and which will continue to play a fundamental role in our present investigations.

The key approach in the present work is the development and testing of formal computational models that account jointly for choice probabilities and complete distributions of *response times* (RTs) in visual WM tasks. To date, researchers in this field have not considered the RT predictions of discrete-slots versus shared-resource models. By collecting and modeling RT data in visual WM tasks, we may open new windows into the underlying processes that are not revealed by accuracy data alone.

Specifically, we develop and evaluate two families of models of visual WM that predict both accuracy and RT. One family formalizes the idea that responding in visual WM tasks emerges from the mixture of cognitive states posited by the discrete-slots view. The second family is based on the assumption of a sharing of continuous resources. Before describing these models, we provide some background on the manner in which discrete-slots models have been applied to the prediction of ROC data.

Formal Modeling of ROC Curves in Visual WM

The starting point for the present investigation is a study by Rouder et al. (2008), who provided intriguing support for the discrete-slots view of visual WM. These researchers used a well-known version of a visual WM task involving change detection (Luck & Vogel, 1997)—see Figure 1A. At study, subjects were presented with an array of colored squares. At test, a single “probe” square was presented at one of the study locations. The observer’s task was to decide whether the color of the probed study square changed or stayed the same. The independent variables that were manipulated were the set size of the study array (2, 5, or 8 study squares) and the *a priori* probability that a “change” occurred in each block of trials (.3, .5, or .7).

Rouder et al. (2008) developed a class of discrete-slots models for predicting performance in the task. In the following, we provide a brief sketch of a representative example from the class. The model assumes that with some probability m_i (which is a function only of set-size i), the observer maintains the relevant study square in WM. If the observer enters this “memory” state, then he or she can judge “change” or “no change” with perfect accuracy (because the colored squares were easily discriminable in Rouder et al.’s, 2008, paradigm). Alternatively, with probability $(1 - m_i)$, the observer fails to maintain the studied square in WM, in which case she is forced to guess whether a change has occurred. The probability that the observer guesses “change,” g_j , is assumed to depend only on the change-probability level (j) on that block of trials. Presumably, the observer tends to guess “change” with higher probability as the *a priori* probability of “change” trials increases.

In the following, a “hit” refers to a case in which an observer correctly detects a change when it occurs, whereas a “false alarm” refers to a case in which an observer guesses “change” when the color of the square remained the same. Let h_{ij} and f_{ij} denote the “hit” and “false alarm” probabilities for trials in which set-size had level i and change-probability had level j . Then according to the representative model sketched above, the hit and false alarm probabilities are given by

$$h_{ij}=m_i+(1-m_i)g_j \quad (1A)$$

$$f_{ij}=(1-m_i)g_j \quad (1B)$$

For example, a “hit” occurs either if the probed study-square location occupies one of the discrete slots of WM (which occurs with probability m_i), or if that study-square has been lost from WM but the observer guesses correctly (with probability g_j) that a change occurred. False alarms arise only if the observer fails to maintain the probed study square in one of the discrete slots and then guesses “change.”¹

Plotting hits against false alarms, it is straightforward to see that the Equation-1 model predicts linear ROC (isosensitivity) curves with slope = 1 and y-intercept m_i (see Figure 1B). As illustrated in the figure, the y-intercept (m_i) of each isosensitivity curve will presumably decrease with increases in memory set size (M). That is, assuming that set sizes are used such that larger sizes increasingly exceed the average number of discrete slots, then the probability that a test item will occupy one of the slots will decrease with the increases in set size. As noted by Rouder et al. (2008, p. 5976), the model also predicts linear isobias curves (i.e., the dashed lines in Figure 1B). Rouder et al.’s (2008) empirical data were well described by these linear ROC curves. Furthermore, specific members from this class of discrete-slots models provided more parsimonious accounts of the data than did alternative signal-detection models based on sharing of continuous resources, and the discrete-slots models were favored by a variety of model-selection statistics.

Outline of New Theoretical Development

Although Rouder et al.’s (2008) results are intriguing, the support for the discrete-slots model was based on limited data. For example, the ROC curve for each memory-set size i was based on only three data points. Although one approach to developing more exacting tests is simply to extend the number and the range of the set-size and change-probability conditions, we propose here instead to test the alternative theoretical views in a fundamentally new way, namely, through the joint prediction of choice probability and RT data.

In this research, we develop two families of models. The first formalizes discrete-slot assumptions, and the second formalizes sharing of continuous resources. We then use the models to predict detailed RT distributions and choice probabilities in a set of visual change-detection tasks. To illustrate, we start by describing a very strong special-case member from the discrete-slots family as applied to Rouder et al.’s (2008) task, and we then outline some natural generalizations.

The basic idea underlying the discrete-slots models is represented schematically in Figure 2. We presume that when a location from the study array is probed, there is an initial “gating”

¹The Equation-1 change-detection model is closely related to the well-known single high-threshold model of signal detection but is not formally identical to that model—for example, see Macmillan and Creelman (1991, pp. 89–90).

process that allows the system to determine whether memory-based information has been retained at that location (i.e., whether the study square is still in WM). Because the central goal of our research is only to determine whether the data are better characterized by discrete-slots or shared-resources views, we do not model the detailed mechanisms that determine the outcome of the gating process. (However, we consider issues related to the gating process in greater detail in our General Discussion.) If the gating process reveals that information has been retained at the probed location, then a memory-based evidence-accumulation process operates that allows the observer to make a “change” or “no-change” decision. Furthermore, for the present version of the paradigm, in which the targets and lures are highly discriminable, the memory-based accumulation process is presumed to operate with extremely high accuracy. Alternatively, if the probed location is no longer in WM, then a “guessing”-accumulation process operates. As indicated in Figure 2, following Rouder et al. (2008), according to this discrete-slots model, the memory-based process operates with probability m_i , which is a function only of set-size i , whereas the guessing process operates with probability $1 - m_i$. To maintain generality, in the present research, the probability m_i is allowed to be a free parameter for each individual set size. In our General Discussion, we sketch approaches to developing more constrained versions from this class of models.

Consider first the case involving change trials. According to a very strong, special-case version of this model, on change trials there is some *single* RT distribution associated with being in the “memory” state, independent of the set size on that trial. The intuition is that the observer either does or does not remember the color of the individually probed square. If she *does* remember, and a change has occurred, then the memory-based process yields a “change” response in accord with that memory-based RT distribution (left branch of Figure 2). Alternatively, the observer may instead enter into the “guess” state (right branch of Figure 2) because the probed location is no longer in WM. Again, following the intuition outlined above, we imagine that there is some single RT distribution associated with guessing “change,” which is independent of study-set size. However, the guessing RT distribution would almost certainly vary with change probability because observers would presumably be more willing to guess “change” on blocks with increased change probability.

According to this strong version of the model, the conditional probability distribution of “hit” RTs as a function of set-size i and change-probability j would be given by

$$f_H(t|i, j) = [m_i \cdot f_{MC}(t) + (1 - m_i) \cdot g_j \cdot f_{GC}(t|j)] / [m_i + (1 - m_i) \cdot g_j] \quad (2)$$

where $f_{MC}(t)$ is the “memory-based” RT distribution associated with correct change responses, and $f_{GC}(t|j)$ is the “guess-change” RT distribution at change-probability level j .²

As described in more detail below, in the family of proposed formal models, the memory-state and guessing-state RT distributions will be generated from explicit evidence-accumulation processes. For now, however, it is useful to consider the descriptive form of

²The numerator of Equation 2 gives the unconditional probability density that, on a change trial, the observer correctly responds “change” at time t . The conditional probability density for the “hit” RTs is then found by dividing this unconditional probability density by the overall probability of a hit.

Equation 2. In particular, this strong version of the discrete-slots model states that, across the set-size conditions i , the “hit” RT distribution is a probabilistic *mixture* of basis distributions $f_{MC}(t)$ and $f_{GC}(t|j)$, with (normalized) mixing parameters m_i and $(1 - m_i) \cdot g_j$. Thus, in addition to quantitative fitting of the detailed RT distributions, we can take advantage of well-known results involving mixture distributions to test fundamental predictions from the model (e.g., Meyer, Yantis, Osman, & Smith, 1985).

A simple example is that the mean RT for hits will be a weighted mixture of the means of the component memory and guessing RT distributions,

$$\mu_{H(i,j)} = [m_i \cdot \mu_{MC} + (1 - m_i) \cdot g_j \cdot \mu_{GC(j)}] / [m_i + (1 - m_i) \cdot g_j] \quad (3)$$

where $\mu_{H(i,j)}$ is the mean of the hit RT distribution, μ_{MC} is the mean of the memory-based change RT distribution, and $\mu_{GC(j)}$ is the mean of the guess-change RT distribution at change-probability level j . A plausible result that one might observe in the RT data is that mean hit RTs will get slower as memory set-size increases. Although such a result is compatible with the Equation-3 model, note that it would not be predicted because of changes in the speed of the memory-based accumulation process—the speed of that process is invariant with set size. Instead, such a pattern would arise if the guessing-based accumulation process operates more slowly than the memory-based one: As set size increases, the mixing probability $(1 - m_i) \cdot g_j$ increases, so there is a greater relative proportion of trials in which the slower guessing process produces the response.

An example of a strong qualitative prediction from the model involves the false-alarm RT distributions. In particular, assuming that the memory-based accumulation process operates with perfect accuracy,³ this strong model predicts that the conditionalized “false-alarm” RT distributions are given by

$$f_{FA}(t|i, j) = f_{GC}(t|j) \quad (4)$$

That is, the observer produces a false-alarm only if she guesses “change” while in the guess state. Thus, from Equation 4, one can see that the conditionalized false-alarm RT distributions are predicted to be *invariant* with memory set size i . Further, because the entire distribution is invariant with set size, the model obviously also predicts that derived summary statistics such as the mean and variance of the false-alarm RTs are invariant with set size.

Although the above reasoning pertained to the patterns of “hit” and “false alarm” RTs, it is straightforward to see that a directly analogous line of reasoning applies to the RTs for correct rejections and misses. In particular, across the different set sizes, the correct-rejection RT distribution will be a probabilistic mixture of a memory-based no-change

³As will be seen, for the best-fitting parameters of the models, this perfect-accuracy condition is a close approximation for most but not all subjects.

distribution and a guessing-based no-change distribution. Furthermore, analogous to the false alarms, the conditionalized “miss” RT distributions will be invariant with set size.

To the extent that such properties are observed in the empirical data, it would constitute strong evidence in favor of the mixture-of-discrete-states view. Continuous shared-resource models have no natural basis for making such *a priori* predictions and will tend to provide poor quantitative fits to RT distributions that satisfy these properties.

The family of mixed-state models that we develop uses evidence-accumulation models to generate the $f_M(t)$ and $f_G(t|j)$ basis distributions. As described below, special cases of these evidence-accumulation models yield $f_M(t)$ and $f_G(t|j)$ distributions that are invariant with memory set size, thereby producing strong models of the form summarized in Equation 2. However, natural generalizations will relax the strong assumption that the “memory-based” RT distribution $f_M(t)$ is completely invariant with set size. For example, in one generalization “drift rate” in the memory-based accumulation process is invariant with set size, but other parameters, such as response thresholds, vary. In these generalizations, the basic spirit of the discrete-slots models is left intact. That is, there are still two basic cognitive states, a memory state and a guessing state. Furthermore, within the memory state, resolution is presumed to be fixed across conditions, as formalized by an invariant drift rate parameter. However, decision processes that operate upon those states may vary with experimental conditions. As will be seen, these generalizations grant the discrete-slots models more flexibility, but their predictions can still be sharply contrasted with those from continuous shared-resource models.

Linear-Ballistic Accumulator Approach

In the following, we outline the general approach to implementing both the discrete-slots models and the continuous shared-resource alternatives. We provide further details regarding specific members of the two families in the Modeling Analyses section. Following recent approaches to modeling short-term memory-scanning phenomena (Donkin & Nosofsky, 2012a, 2012b; see also Nosofsky, Little, Donkin, & Fific, 2011), we implement the visual change-detection models by using a linear ballistic accumulator (LBA) approach (Brown & Heathcote, 2008). The LBA model is computationally simple and allows analytic computation of trial-by-trial likelihoods. Furthermore, recent research indicates that the LBA and the extremely influential diffusion model (e.g., Ratcliff, Van Zandt, & McKoon, 1999) are able to mimic one another’s predictions closely and that parameters that share common psychological interpretation map closely onto each other (Donkin, Brown, Heathcote, & Wagenmakers, 2011).

To illustrate our LBA-based modeling approach, first consider the memory-based accumulation process. According to the LBA model, there are two accumulators, one for change responses and the second for no-change responses (Figure 3). Evidence accumulates in linear and ballistic fashion (i.e., with no within-trial variability in evidence accumulation) until a response threshold is reached on either accumulator. Evidence accumulates with mean drift rate v_C on the change accumulator and with mean drift rate v_{NC} on the no-change accumulator. (If a stimulus change actually occurred, then v_C would tend to be large in

magnitude and v_{NC} would tend to be small; and vice-versa if no stimulus change occurred.) If the accumulating evidence first reaches the “change” response threshold (R_C), then the observer responds “change”; whereas if the evidence first reaches the “no change” response threshold (R_{NC}) then the observer responds “no change.” The decision time is determined by the time that it takes the accumulating evidence to reach either threshold. Different versions of the model arise by varying assumptions about how drift rates and response thresholds may vary with experimental conditions (see Modeling Analyses section).

An analogous scheme operates for the guessing-based accumulation process, except here drift rates are not stimulus-dependent. Instead, we assume some constant mean drift rate v_G in both a guess-change and guess-no-change accumulator. (Because the guessing drift rate is not stimulus-driven, we refer to it generically as an “accumulation” process rather than as an “evidence”-accumulation process.) The observer is presumed to vary the response thresholds on these guessing accumulators in accord with experimental conditions. For example, in conditions in which change-probability is high, the observer would presumably set a low threshold on the “guess-change” accumulator, and a higher threshold on the “guess-no-change” accumulator, leading to more probable and faster “guess-change” responses.

In all of the discrete-slots models, under the present experimental conditions, a fundamental assumption is that the mean drift rates v_C and v_{NC} on the memory-based accumulators are *independent* of memory set size and change probability. This assumption formalizes the “all-or-none” conception that is integral to the discrete-slots view. In other words, because drift rate in the model is determined by memory resolution, and resolution is presumed to be invariant with set size in the discrete-slots models, then drift rate is invariant with set size. The key parameters that vary across the different memory-set size conditions (for *all* versions of the discrete-slots models) are the memory-state probability parameters, m_i (Figure 2). As explained in detail in the Modeling Analyses section, different versions of the discrete-slots models are produced by making alternative assumptions about how response thresholds on the memory-based accumulators vary with experimental conditions.

Finally, it is straightforward to use this LBA framework to develop competing models that formalize the continuous shared-resource view of visual WM. Instead of assuming that there is a mixture of cognitive states (memory plus guessing), there is now only a memory-based evidence-accumulation process (i.e., the left branch of Figure 2). However, recall that for the discrete-slots models this memory-based accumulation process is highly constrained: Regardless of memory-set size and change probability, the mean drift rate on the change accumulator is given by an invariant v_C . By contrast, for the continuous shared-resource models, the drift rates v_C are allowed to vary with experimental conditions. Presumably, for example, as memory set-size increases, mean correct drift rates on the change accumulator will decrease because the observer has stored lower-resolution representations of the to-be-remembered items. (In Experiment 2, we will also consider hybrid models of visual WM change detection, which assume both a mixture of cognitive states involving memory and guessing, in addition to memory-based drift rates that vary with set size.) Analogous to the discrete-slots models, a variety of versions of the continuous shared-resource models can be formalized, depending on precise assumptions that are made about how drift rates and

response thresholds vary with experimental conditions (see Modeling Analyses section for details).

It is important to note that there is a sharp conceptual and mathematical contrast between the discrete-slots and continuous shared-resource families. It is true that, at a gross level of analysis, both models predict that performance will tend to decline as set size increases. In the discrete-slots models, the main reason for the decline is that the probability of entering the memory state (m_i) decreases as set size increases. In the continuous models, the main reason for the decline is that correct drift rates (v) decrease as set size increases. But the two families of models are not simply exchanging one parameter for another. The discrete-slots models posit that performance arises from a probabilistic mixture of memory-based and guessing-based processes, whereas the shared-resource models posit continuous variation in a single memory-based evidence-accumulation process. The types of RT distributions and ROC curves that are produced by the mixture process can be very different in form from those produced by the continuous models. Furthermore, as we described earlier, only the discrete-slots models make the *a priori* qualitative prediction that the error RT distributions should be nearly invariant with set size. As will be seen, in many cases, these contrasts are sufficient to sharply discriminate between the quantitative predictions from the two competing families.

Experiment 1

The purpose of this experiment was to collect choice-probability and RT data in a visual WM change-detection task suitable for contrasting models from the discrete-slots and shared-resources families. To achieve this goal, we conducted a partial replication and extension of the paradigm reported by Rouder et al. (2008). The design was essentially the same as the one reported by Rouder et al. (2008). However, because the goal was to model detailed RT-distribution data at the individual-subject level, each observer contributed more trials of data than in the original experiment.

Method

Participants—Eight participants from the Indiana University community were paid \$12 per session (which included a \$3 bonus for good performance) to complete 10 sessions of the task. All participants reported having normal color vision, and none was aware of the issues under investigation in the study.

Stimuli—Stimuli were chosen from a set of 10 dissimilar color squares (white, black, red, blue, green, yellow, orange, cyan, purple, and dark-blue-green). Stimuli were presented in the same way as was described by Cowan et al. (2005): Items were presented within an array whose visual angle was within a 9.8×7.3 degree rectangle. Each color square was 0.75×0.75 degrees in size. Items were presented in randomly chosen locations with the constraint that they had to be at least 2° away from any other item and from the center of the viewing area. The cue indicating the position of the probe color square was a black, 1 pixel thick, 1.5 degree diameter circle that surrounded the probe color square. Stimuli were presented on a grey background on 17" cathode ray tube monitors.

Procedure—The structure of trials is outlined in Figure 1A. A fixation cross was presented for 1,000 ms to begin the trial. A study array of N color squares was then presented for 500 ms. After a 500 ms pause, a multicolored pattern mask was presented for 500 ms. Following the mask, a single test color was presented in one of the locations of the study array. The test color was either the same as the item in that location of the study array (a “no-change” trial) or was different from the item in that particular location of the study array (a “change” trial). The participant was asked to indicate whether the test item was the same as the study item or had changed, by pressing “F” or “J” on the keyboard, respectively. After a 1,000 ms break, the next trial began.

For purposes of generality, we decided to test both an “internal”-change and an “external”-change condition, with four of the eight subjects participating in each condition. In the internal condition, the test item on change trials was an item from the study array not in the test location. That is, the test item changed into another item that had previously been presented in the study array. In the external condition, the test item on change trials was not an item presented in the study array. As it turned out, the internal versus external manipulation had little effect on the main pattern of results in this experiment, so we collapse across this variable in our ensuing presentation. As described in some detail in our General Discussion, however, the internal versus external distinction is still of potential importance in understanding the details of human performance in this task.

In the external condition, the number of items in the study array, that is, the set size N , was 2, 5, or 8. In the internal condition, N was 3, 5, or 8.⁴ Set size varied randomly from trial-to-trial, with the constraint that each N was presented 20 times in each block of 60 trials. The proportion of change trials was also manipulated: In a given block, the test item differed from the study item on either 30%, 50%, or 70% of trials. The number of change (and no-change) trials was equal across all values of N within each block. In each session, participants completed three blocks of each change probability condition, for a total of nine blocks (or 540 total trials). Every three blocks of trials, each of the change-probability conditions was tested once. The order of testing within each of the three blocks was randomized. Following Rouder et al. (2008), we informed participants of the change-probability manipulation at the start of each block of trials. This information was provided by both telling the participant the percentage of upcoming trials on which the test item would change and by using a pie chart to illustrate the manipulation.

Results

We deleted from the analysis any trial in which the RT was less than 180 ms or greater than 2,500 ms. In addition, for each combination of set size, change probability, and response type (i.e., hit, miss, false alarm, and correct rejection), we deleted from the analysis any trial in which the RT was greater than three standard deviations above the mean RT for that combination. These procedures led to deleting at most 3.4% of the trials for any subject. In these trimmed data, summary results based on medians were qualitatively similar to those based on means, so only the means are displayed.

⁴We conducted the internal condition after the completion of the external condition. Inspection of the external-condition data revealed essentially no errors for $N = 2$, so we decided to increase the smallest set size to $N = 3$ in the internal condition.

Although our main goal involves formal modeling of the individual-subject data, we start by providing a brief description of the general pattern of results. The choice-probability and mean RT data, averaged across the eight subjects, are displayed in Figure 4. The patterns of averaged data displayed in the figure are reasonably representative of the patterns that we observed at the individual-subject level. We will describe some individual differences, however, in the Modeling Analyses section below.

As can be seen in the top-left panel of Figure 4, as set size increased, hit rates decreased while false-alarm rates increased. That is, performance grew less accurate with increases in memory set size. The decrease in hit rates across set-size conditions was statistically significant [$F(2, 14) = 45.9, p < .001$], as was the increase in false-alarm rates [$F(2, 14) = 38, p < .001$]. In addition, as shown in the top-right panel, mean hit rates and mean false alarm rates both increased with increases in change probability. That is, subjects tended to guess “change” more often as change-probability increased. These increases were statistically significant for both hits [$F(2, 14) = 12.4, p = .001$] and false alarms [$F(2, 14) = 18.6, p = .001$]. These qualitative effects of set size and change probability on the hit and false-alarm probabilities are naturally predicted by all of the models.

To investigate the joint effects of change probability and set size on the choice probabilities, we inspected the ROC curves of the individual subjects. At the individual-subject level, the ROC plots tended to be noisy. Furthermore, although the effect of change probability on the hit and false-alarm rates was statistically significant, the magnitude of the effect tended to be small, so the three points that defined each curve tended to be close together. Taken together, these factors made it difficult to determine whether the individual-subject ROC plots were straight lines (as predicted by the discrete-slots models) or curved (as predicted by the continuous models). The averaged ROC data appeared to be curved, but the three points that defined each curve were again close together. (Rouder et al., 2008, observed a similar result, despite the fact that their individual-subject modeling tended to favor the assumption of straight-line ROCs.) However, the shape of averaged ROC curves needs to be interpreted with a good deal of caution. For example, if the performance of most observers follows the discrete-slots models, but a subset obeys the continuous models, then the averaged plot would be curved. Because the present models are required to simultaneously fit both the RT-distribution *and* choice-probability data of the individual subjects, the structure of the individual-subject ROCs places important constraints on the competing models.

The mean RT data are displayed in Rows 2–4 of Figure 4. Consider first the correct RTs (i.e., the hit and correct-rejection RTs), which are displayed as a joint function of set size and change probability in Row 2. First, both hit and correction-rejection RTs got slower as set size increased. Second, as change probability increased, hit RTs got faster, whereas correct-rejection RTs got slower. To confirm these observations, we conducted a two-way ANOVA on the mean RT data using set size (2/3, 5, 8) and change probability (.3, .5, .7) as factors. For hit RTs, there was a main effect of set size [$F(2, 14) = 27.3, p < .001$], a main effect of change probability [$F(2, 14) = 11.4, p < .001$], and no interaction ($p < .22$). Likewise, for correct rejections, there was a main effect of set size [$F(2, 14) = 62.1, p < .$

001], a main effect of change probability [$F(2, 14) = 58.2, p < .001$], and no interaction ($p < .90$).

The pattern of results for set size is naturally accounted for by continuous shared-resource models, assuming that correct drift rates get slower with increases in set size. However, the pattern is also accounted for by discrete-slots models for a variety of reasons. For example, if the guessing process tends to be slower than the memory-based process, then there will be a slow-down with set size because a greater proportion of the responses will be based on guessing as set size increases. As will be seen, however, the magnitude of the set-size effects will place important constraints on specific models from the two main families. Likewise, assuming that observers adjust their response thresholds to guess “change” more readily under conditions in which change probability increases, then the effects that change-probability has on the hit and correct-rejection RTs are also naturally predicted by both families of models. Once again, however, it is an open question whether specific members of the two families of models can predict the magnitude of the effects.

The corresponding mean RTs for false alarms and misses are displayed in Figure 4 as a function of change-probability (Row 3) and set size (Row 4). Because these error RTs are based on much smaller sample sizes compared to the hit and correct-rejection RTs, we display the results separately for each independent variable. As expected, miss RTs got slower and false-alarm RTs got faster with increases in change probability (Row 3), reflecting subjects’ increased tendency to respond “change.” Because of the increased variability in the data, these effects failed to reach statistical significance in the group averages ($ps > .24$); however, as will be seen, some individual subjects showed pronounced effects of change-probability on the error RTs.

The set-size effects on the error RTs are shown in Row 4. The results for the smallest set size need to be interpreted with a good deal of caution because they are based on very low frequencies. (Subjects rarely made errors for the smallest set size—see top-left panel of Figure 4.) To emphasize this fact, we have plotted the error RTs for the smallest set size with open rather than solid symbols. Nevertheless, the right panel of Row 4 shows that set size had virtually no effect on miss RTs. This interesting result is in accord with the *a priori* predictions from the discrete-slots models. For the miss RTs, the effect of set size did not approach statistical significance ($F < 1, p = .98$). The results for the false-alarm RTs are less clear. On the one hand, the mean RTs for set-size 5 and set-size 8 were nearly equal, and the difference was not statistically significant ($p = .11$). (The small apparent difference between set sizes 5 and 8 was due to the behavior of one or two subjects.) On the other hand, despite the very small sample sizes that are involved, the mean false-alarm RT for the smallest set size was significantly faster than for sizes 5 and 8 [$F(2, 14) = 13.5, p < .001$], and the magnitude of the effect was large. Moreover, our examination of the individual-subject data revealed that this pattern was quite consistent across almost all of the subjects. We discuss the result in more detail after presentation of the formal-modeling analyses and suggest in our General Discussion that it may arise from rare trials in which the observer fails to attend or in which feature migrations occur. Crucially, because error RTs at the smallest set size are based on an extremely small number of observations, even rare contaminating events can have an extremely large influence on the computed mean RT.

Modeling Analyses

Parameterizations of the formal models—In this section, we describe the parameterizations that instantiate specific members of each family of visual WM models. We provide a glossary of the parameters of the alternative models in Table 1.

Discrete-slots models: We start by describing the parameters that are common to all of the discrete-slots models. All models from the discrete-slots family estimate three memory-state probabilities, m_i ($i = 1-3$), representing the probability that the item in the probed location is stored in one of the discrete slots when set size is at level i (left panel of Figure 2). (The probability that the observer instead needs to guess is given by $1 - m_i$.) All models from the discrete-slots family also estimate a single mean correct drift-rate parameter v_C for change trials and a single mean correct drift-rate parameter v_{NC} for no-change trials (Figure 3). (Following past applications of LBA approaches, the mean incorrect drift rates on the competing accumulator on these types of trials are given by $1 - v_C$ and $1 - v_{NC}$, respectively.) All discrete-slots models incorporate a single mean drift rate v_G on the guess-change and guess-no-change accumulators. Differential guessing tendencies are modeled across conditions by allowing the response thresholds on the guess-change and guess-no-change accumulators to vary as a function of response type and change probability. Thus, there are six guessing-related response-threshold parameters, $G_C(j)$ and $G_{NC}(j)$ [$j = 1-3$], where the index j denotes level of change probability. In LBA modeling, the start-point of the evidence accumulation process varies randomly across trials in accord with a uniform distribution (see Figure 3). For the memory-based accumulators, the range of the uniform distribution is $[0, A_M]$, and for the guessing accumulators, the range is $[0, A_G]$. Other parameters that are common to all the discrete-slots models are the between-trial (normally distributed) standard deviation of drift rate on the memory accumulators (σ_M) and the guessing accumulators (σ_G) as well as a residual base time t_0 that is not associated with change-detection decision making. (The standard deviation of the guessing drift rate, σ_G , scales the values of the other LBA-related parameters, and without loss of generality we hold it fixed at $\sigma_G = .10$.)

The strongest (i.e., most constrained) model from the discrete-slots family estimates a single response-threshold parameter for the memory-based change accumulator (R_C) and a single response-threshold parameter for the memory-based no-change accumulator (R_{NC}). This strong version (denoted model DS1) incorporates a total of 18 free parameters (see Table 1 for a listing). Assuming that the memory-based accumulation process yields perfectly accurate responding (which, for most subjects, turns out to be a fairly close approximation for the best-fitting parameters from the model), then this strong version yields predictions that are captured by the descriptive Equation 2.

In the first generalization (denoted model DS2), the response thresholds on each of the memory-based accumulators are allowed to vary as a function of memory set size, thereby adding four free parameters to the strong model (for a total of 22 free parameters). In modeling short-term memory-scanning data, various researchers have obtained evidence that observers adjust response thresholds with changes in set size (e.g., Donkin & Nosofsky, 2012b; McElree & Doshier, 1989; Nosofsky et al., 2011; Ratcliff, 1978). To the extent that

similar principles operate in this domain of visual change detection, this generalization of the model is likely to be an important one. Note that the predictions from this generalized model are no longer captured by Equation 2 because the $f_M(t)$ distribution is no longer invariant across the different set-size conditions. Interestingly, however, this generalized model still predicts that the conditionalized false-alarm RT distributions (as well as conditionalized miss RT distributions) should be nearly invariant with set size (see Equation 3). The reason is that false alarms and misses are almost always generated via the guessing accumulators.

A further generalization (DS3) allows the response thresholds to vary jointly as a function of set size *and* change probability. For simplicity, we assumed that the contributions of set size and change probability to the response-threshold settings were additive. So, for example, the response threshold setting on the memory-based change accumulator at set-size (*ss*) level *i* and change-probability (*cp*) level *j* would be given by $R_C(i,j) = ss_C(i) + cp_C(j)$ and analogously for the response-threshold setting on the no-change accumulator. Generalizing the model in this manner adds another 4 free parameters, yielding a total of 26 free parameters in the DS3 model.

Continuous shared-resource models: For the continuous shared-resource models, there is only a single pair of memory-based accumulators, one that accumulates evidence toward a “change” response and the other toward a “no-change” response. However, in accord with the shared-resources idea, drift rates on these accumulators are presumed to vary continuously with memory set size. The baseline version of the continuous model defines three separate mean correct drift rates on the change accumulator and three separate mean correct drift rates on the no-change accumulator. We denote the drift rates $v_C(i)$ and $v_{NC}(i)$, where, for example, $v_C(i)$ denotes the mean correct “change” drift rate at set-size-level *i* ($i = 1-3$). (Again, following previous LBA applications, the mean drift rates on the incorrect accumulators are given by one minus the mean correct drift rates.) In addition, the baseline model defines three separate response-threshold parameters on the change accumulator [$R_C(j), j = 1-3$] and three separate response-threshold parameters on the no-change accumulator [$R_{NC}(j), j = 1-3$], where “*j*” refers to the different levels of change probability. Finally, analogous to the discrete-slots models, there is a start-point variability parameter A , a drift-rate variability parameter σ , and a residual base-time parameter t_0 . This baseline model from the continuous family (denoted model C1) uses 15 free parameters.

The first generalization (model C2) of the baseline model allows the response thresholds to vary as a function of set size as well as change probability. As is the case for the most general of the discrete-slots models, the response-threshold setting on the change accumulator at set-size level *i* and change-probability level *j* is given by $R_C(i,j) = ss_C(i) + cp_C(j)$ and likewise for the response-threshold setting on the no-change accumulator. This generalization adds 4 free parameters to the baseline model, yielding a total of 19 free parameters. Finally, for reasons that are made clear below, the most general of the continuous shared-resource models estimates mean *error* drift rates on each accumulator separately from the correct drift rates. That is, this most general model does not obey the constraint that the error drift rate is given by one minus the correct drift rate. With an additional six mean error drift-rate parameters [$v'_C(i)$ and $v'_{NC}(i), i = 1,3$], this

generalization of the continuous model (C3) uses a total of 24 free parameters. (In model C3, the drift-rate variability parameter σ can be set at arbitrary positive value without loss of generality.)

Model-fitting method—We fitted the models to the individual-subject data by conducting computer searches for the values of the free parameters that yielded maximum likelihood fits to the *individual-trials* choices and RTs. In other words, in this model-fitting method, the likelihood of the choice and RT on each individual trial is assessed, and the overall likelihood is the joint likelihood of all of the individual-trial likelihoods. The analytic equations for expressing these likelihood-based fits of the standard LBA model are presented by Brown and Heathcote (2008, Equations 1–3), and it is straightforward to extend those equations for fitting the mixed-state (i.e., discrete-slots) LBA models. The fit of each model was then assessed by using the Bayesian information criterion (BIC; Schwarz, 1978),

$$\text{BIC} = -2\ln(L) + p\ln(n),$$

where L is the (maximum) likelihood, p is the number of free parameters in the model, and n is the total number of trials on which the fit is based. (Averaged across subjects, the number of data points, n , was equal to 5,180.) The term $p\ln(n)$ is a penalty term that penalizes a model for its number of free parameters. The model that yields the smallest BIC is considered to provide the most parsimonious account of the data. In an attempt to avoid local minima, we used 40 different random starting configurations of the parameters in conducting the computer searches.

We should emphasize that the goals for the modeling are quite ambitious: There are 9 conditions (3 set-size by 3 change-probability conditions, fully crossed), and for each condition we are attempting to predict RT distributions associated with hits, misses, false alarms, and correct rejections (as well as the probability of each of those responses). Thus, the models are required to fit, at the level of individual subjects, individual-trials data that compose 36 distinct unconditional RT distributions.

Model-fitting results—The model-fitting results are reported in Table 2. The fits are reported in terms of BIC values, which are calculated by taking the difference between each model's BIC and the model with the smallest BIC value for that subject. Therefore, a BIC value of 0 represents the best-fitting model for that subject. The greater the BIC value, the worse is the fit of the model for that subject.

Inspection of Table 2 reveals that for seven of the eight subjects, a member from the class of discrete-slots models provided the best BIC fit. In most cases, the improvement in BIC yielded by a member of the discrete-slots family (compared to the continuous family) was quite dramatic. Further, in the single case in which the continuous family achieved a better fit (Subject 7), the improvement in BIC was relatively small. Although the model that yielded the best BIC was usually the one with the greatest number of free parameters (i.e., model DS3), the intermediate discrete-slots model (DS2) also tended to perform quite well.

For example, for six of the eight subjects, model DS2 yielded a better BIC fit than did *any* member of the continuous family.

To confirm the diagnosticity of these model-fitting results, we conducted a series of model-recovery analyses. For each individual subject, we used continuous-model 2 (C2) to generate 20 sets of simulated data. The parameters used to generate the simulated data were those that had provided the best fit of model C2 to that subject's actual data. Each simulated data set was based on the same number of observations in each condition as occurred in the actual experiment. We then fitted model C2 and each of the discrete-state models to each simulated data set, using BIC as the criterion of fit. The results of this model-recovery analysis were overwhelming: In all 160 cases (eight subjects by 20 simulated data sets), model C2 yielded a better BIC fit than did any of models DS1, DS2, or DS3. Thus, it appears that if model C2 were the generating model of the actual data, then use of the BIC statistic would have led to its recovery.

Within the discrete-slots family, there was some variability across individual subjects in the best-fitting model. For five of the eight subjects, the most general version (DS3) provided the best BIC fit, suggesting that those observers adjusted their response thresholds on the memory-based accumulators in accord with both set size and change probability. However, the data for three subjects were better fit by stronger (more constrained) versions of the model. Interestingly, although the strongest member of the family (i.e., DS1, which assumes invariant memory-based RT distributions across all conditions) generally yielded poor fits, a single subject (S5) did appear to behave in accord with its predictions.

To gain insights into these patterns of results, we now consider the model predictions for two individual subjects in some detail. These in-depth inspections are intended to provide deeper insights into the operating principles of the models and to explain the reason why the discrete-slots models tend to be favored. At the same time, in view of the important individual differences revealed by inspection of Table 2, we chose one subject to be representative of model DS2 and a second to be representative of model DS3.

Representative individual-subject predictions—Although the models were evaluated in terms of their fits to 36 sets of RT distributions, our initial focus in this section will be on the patterns of mean RTs and choice probabilities. Recall that subjects rarely made errors at the smallest set size. Thus, the mean error RTs at the smallest set size were based on extremely limited sample sizes and tended to show great variability across individual subjects. Also, because of the extremely small sample sizes, the BIC comparisons were minimally affected by the error RTs in those conditions. Therefore, we do not display those results in the individual-subject figures. We will suggest in our General Discussion that the error responses at the smallest set size likely reflect special processes that go outside the scope of the present models.

Subject 3: The results for Subject 3 are displayed in Figures 5 and 6. Figure 5 shows the predictions from the intermediate discrete-slots model (DS2), which yielded the best BIC fit for this subject. For purposes of comparison, we display the results from the intermediate

continuous model (C2) in Figure 6. To understand more fully the basis for the predictions, in Table 3 we report the best-fitting parameters from the favored model (DS2).

As shown in Figure 5, model DS2 appears to do a reasonably good job of predicting almost all of the major trends for Subject 3. First, consider the accuracy data plotted in Row 1. The model predicts the decrease in accuracy (decrease in hits and increase in false alarms) as a function of set size (top left panel) and the increase in probability of responding change (both hits and false alarms) as a function of change probability (top right panel). It predicts the set-size effects on accuracy because, as set size increases, the probability of entering the memory state (m_i) decreases, so the observer must rely on the guessing process. It predicts the effects of change-probability on hits and false alarms because, as change probability increases, the observer adjusts the response thresholds on the guessing accumulators to yield more “change” responses. In particular, the observer sets a more lenient response threshold on the guess-change accumulator and a stricter response threshold on the guess-no-change one (see Table 3).

Next, consider the RT data, plotted in Rows 2–4. The model accurately predicts the slow-down in mean correct RTs (hits and correct rejections) with increases in set size (Row 2). It predicts these results because, as set size increases, there is an increased probability of entering the guessing state, and the guessing process operates more slowly than does the memory-accumulation process for Subject 3. In addition, the subject sets more conservative response thresholds on the memory-based accumulators as set size increases (Table 3). The model also predicts that, in general, as change probability increases, hits get faster and correct rejections get slower (Row 2). It predicts these results because, as discussed above, the observer sets a more lenient threshold on the guess-change accumulator (and a stricter threshold on the guess-no-change accumulator) as change probability increases. However, the model underestimates the magnitude of the change-probability effects for correct rejections, especially at the largest set size.⁵

Finally, with the exception of false alarms at the smallest set size (not shown), the model also does a good job of predicting the mean error RTs. False-alarm and miss RTs are predicted to vary with change probability (Figure 5, Row 3) because of the adjustments to the guessing response thresholds described above. And, as explained in the introduction, the model predicts correctly the nearly flat set-size effects (Row 4) on the miss RTs and the false-alarm RTs (for set sizes 5 and 8). These error responses are generated almost entirely via the guessing accumulators, whose parameters are invariant with set size.

Recall that the highly constrained discrete-slots model (DS1) performed poorly for almost all subjects. Although not illustrated graphically, the main reason for the poor fits of model DS1 is that it failed to predict the magnitude of the slow-down in the hit and correct-rejection RTs as set size increased. Recall that the baseline DS1 model can predict slower RTs for correct responses with increases in set size only because of an increased proportion of guessing responses (see mixture- Equation 3). Thus, the degree of predicted slow-down for the mean correct RTs is tightly constrained by the mean RT for the guesses that result in

⁵This type of misprediction was even more pronounced for some of the other subjects, which necessitated the use of model DS3.

error responses (i.e., the false alarms and misses). The observed magnitude of the slow-down, however, was greater than predicted by this mechanism alone. This same failing of the strong version of the discrete-slots model held for seven of the eight subjects (the one exception was Subject 5).

For purposes of comparison, the predictions from the continuous model (C2) are displayed in Figure 6. Perhaps the main limitation of that model is that it appears to mispredict many of the error RTs. Because the correct RTs are based on much larger sample sizes, the fit routine probably settles on parameter values that allow the continuous model to accurately predict that subset of the data. But the model cannot simultaneously handle both correct and error RTs. One reason why the model has trouble involves assumptions about how error drift rates are related to correct drift rates. Recall that in standard LBA applications, the mean drift rate on the error accumulator is set at one minus the mean drift rate on the correct accumulator. The main mechanism by which the continuous model predicts a slowing of RTs for correct responses with increases in set size is to assume that mean correct drift rates decrease with set size (due to the storage of lower-resolution memory traces). A consequence of this assumption is that incorrect drift rates *increase* with set size, leading to the prediction of *faster* errors as set size increases. The continuous model attempts to “undo” this incorrect prediction by increasing the magnitude of the response thresholds as set size increases. But it is only partially successful in remedying the difficulty. For example, for Subject 3, the model predicts that false-alarms will get slightly *faster* as one moves from set-size 5 to set-size 8 (Figure 6, lower-left panel). It is because of this difficulty that we tested the more elaborate version of the continuous model in which error drift rates were allowed to vary freely as a function of set size. However, the continuous model with freely estimated error drift rates still failed to yield BIC fits that were as good as those yielded by members from the discrete-slots family. Perhaps the reason is that the model has no natural basis for predicting *a priori* that the distribution of error RTs will be nearly invariant across the different set-size conditions.

Subject 1: In our second example, we consider results from Subject 1, whose data were best fit by the most complex of the discrete-slots models (DS3). Figures 7 and 8 show the predictions from models DS3 and C2, respectively.

As can be seen in Figure 7, model DS3 accounts well for almost all of the effects of set size and change probability on the choice probabilities and mean RTs. (Although not shown in the figure, it failed to predict a fast false-alarm mean RT at set-size 2.) The reason why model DS3 provided such a better fit compared to model DS2 involved the patterns of correct RTs. In particular, model DS2 failed to capture the extent to which mean hit RTs were slowed in the lowest change-probability condition, as well as the extent to which mean correct-rejection RTs were slowed in the highest change-probability condition. According to model DS2, any effects of change probability on the speed of correct responses are due solely to changes in the speed of the guessing process. Subject 1, however, showed extremely accurate responding in all conditions, so the guessing process does not make a large contribution to the subject’s correct RTs. Instead, by also allowing adjustments in response thresholds on the memory-based accumulators with manipulations of change probability, model DS3 accounts for these effects in a natural fashion. The limitations of the

continuous model (C2) are evident from inspection of Figure 8. Among the main limitations is that it badly mispredicts the speed of the false-alarm RTs at almost all levels of set size and change probability.

RT-distribution predictions—To provide further documentation of the ability of the discrete-slots models to account for the individual-subject data, we provide displays of the observed and predicted RT distributions in Figures 9 and 10. In these figures, the histograms show the empirical RT distributions, whereas the smoothed curves are the predicted distributions from the discrete-slots models. Figure 9 displays the results for Subject 1 (using model DS3), and Figure 10 displays the results for Subject 3 (using model DS2). Each row of these figures shows the RT distributions for a different combination of change probability and set size. The top three rows correspond to change-probability .3; the middle three rows correspond to change-probability .5; and the bottom three rows correspond to change-probability .7. Within each set of three rows, set size increases from 2 to 5 to 8. Note that the area underneath each histogram is proportional to the absolute frequency with which that response type occurred at each change-probability/set-size combination. For example, hits were far more common under high change-probability conditions, whereas correct rejections were far more common under low change-probability conditions.

Inspection of Figures 9 and 10 indicates that it is not only the central tendencies of the RT distributions that change across experimental conditions. Instead, there are also systematic changes in the spreads, shapes, and leading edges of the distributions. For example, consider the hit RT distributions for Subject 1 (left column of Figure 9). Within each change-probability condition, as set-size increases, the distributions change from being narrow and peaked to wider and more positively skewed. Furthermore, the leading edge of each distribution shifts slightly to the right. In addition, as change probability increases (i.e., as one moves from Rows 1–3 to Rows 4–6 to Rows 7–9), the leading edges of the hit RT distributions shift further to the left. The discrete-slots model does a good job of capturing all of these patterns. Analogous patterns are observed for the correct-rejection RT distributions, and these are also well captured by the model. The patterns of observed data for Subject 3 are similar to those of Subject 1. Although the restricted version of the discrete-slots model does not appear to capture the RT distributions with the same precision as those of Subject 1, the overall fits still seem reasonably good.

As noted in our introduction, the discrete-slots models predict that the conditional RT distributions for false alarms and misses should be invariant across the different set size conditions. We have already acknowledged that this strong prediction is disconfirmed for the smallest set size, even at the level of mean RTs. However, it is of interest to test the prediction at the larger set sizes (5 and 8). To do so, we conducted Kolmogorov-Smirnov (K-S) tests for the equality of the RT distributions at set sizes 5 and 8 for each of the individual subjects. In one set of tests, we compared these RT distributions separately at each individual change-probability condition. In the second set, we compared the RT distributions pooled across the change-probability conditions. For completeness, in addition to comparing the error RT distributions (i.e., false alarms and misses), we compared the correct RT distributions (i.e., hits and correct rejections) as well. The *p* values associated with these tests are reported in Table 4. Even without correcting for multiple tests, it can be

seen that the null hypothesis of no difference between the size-5 and size-8 RT distributions is rarely rejected for the false-alarms and misses (see top panel of Table 4). This null result is obtained regardless of whether the analysis is conducted separately at each individual change-probability condition or pooled across conditions. The main exception is for Subject 7, who was also the single subject whose data were better fit by the continuous models than by the discrete-slots models. By comparison, when applied to the hits and correct rejections (bottom panel of Table 4), the K–S tests lead one to reject the null hypothesis of no difference between the size-5 and size-8 RT distributions. The vast majority of tests conducted at the level of individual change-probability conditions are statistically significant, and the results from the pooled analyses are over-whelming. Taken together, the results tend to support the discrete-slots models' predictions that hit and correct-rejection RT distributions will differ across the set-size conditions but that the false-alarm and miss RT distributions will be nearly invariant.

Averaged-subject predictions—In the previous section, we examined the modeling results from only two of the subjects, albeit in considerable detail. However, the modeling results from those two subjects are representative of the larger group. To provide some documentation of this claim, in Figure 4 we display, along with the observed averaged data, the averaged predictions from model DS3. As is apparent from inspection, the model captures in quantitative detail the manner in which the averaged hit and false alarm rates vary with set size and change probability (see Row 1 of Figure 4). It also accounts accurately for how the mean hit and correct-rejection RTs change with these variables (Row 2 of Figure 4). The model captures as well the manner in which false-alarm and miss RTs vary with change probability (Row 3) and captures the flat set-size function associated with miss RTs (Row 4, right panel). Its main limitation, as we acknowledged earlier, is that it fails to capture the fast false-alarm RT at the smallest set size. We consider these results more fully in the General Discussion. Before turning to that broader discussion, however, we first provide further tests of the competing models in an extended version of the change-detection experiment.

Experiment 2

A limitation of Experiment 1 is that the task required the observer to detect only big, easily discriminable changes. In Experiment 2, we extend the design by also including small-change trials. This extension is important for several reasons. First, it is possible that observers have flexible modes of responding in these tasks and that they go into something akin to a “low-resolution mode” when the changes always cross major category boundaries. Will evidence for a mixture of discrete states still be obtained under conditions in which observers must detect small changes as well? Second, one's ability to measure any fine-grained changes in resolution may be limited when only big-change trials are included in the design. By including small-change trials as well, we may find evidence for hybrid accounts involving both discrete-state processing (i.e., memory-based and guessing-based responding) *and* continuous changes in resolution as memory set-size increases.

To investigate these possibilities and provide further theoretical tests, we expanded the stimulus set that had been used in Experiment 1. For each of the original colors in the

stimulus set, we included a new color from the same hue category, but that differed slightly from the original in its brightness and saturation. Thus, across hues, color changes were big, but within hues, color changes were small. The design was the same as in Experiment 1, except that on change trials, half the changes were big and half the changes were small.

It should be noted that some of the qualitative predictions that helped distinguish model classes in Experiment 1 are missing in Experiment 2. In particular, although the discrete-slots models still predict that the distribution of *miss* RTs on big-change trials should be nearly invariant with set size, they no longer predict invariant *false-alarm* RT distributions on *same* trials. The reason is that because of the difficulty involved in discriminating between *same* and *small-change* trials, false alarms on *same* trials will not arise solely from the guessing process but from the memory-based accumulation process as well.

As already noted, the design of Experiment 2 should be helpful in estimating any fine-grained changes in memory resolution with changes in set size and allows one to investigate the utility of *hybrid* models of visual change detection. Another advantage of the new design is that it allowed us to explore the counterintuitive prediction from the discrete-slots models that certain types of responding might get *faster* as memory set size increases. In Experiment 1, the main reason why the discrete-slots models predicted slower *hit* and *correct-rejection* RTs with increases in set size is that there would be a greater proportion of guessing trials, and the guessing process was slower than the memory-based process. However, in the current experiment, discriminating between *same* versus *small-change* trials is difficult, so the memory-based accumulation process related to such decisions may be slow. Under such conditions, responses based on guessing may be *faster* than responses based on memory, leading to the possibility that certain types of RTs may get faster as set size increases.

Method

Participants—Five new participants from the Indiana University community were paid \$12 per session (which included a \$3 bonus for good performance) to complete 10 sessions of the task. All participants reported having normal color vision, and none was aware of the issues under investigation in the study.

Stimuli—The same 10 color squares were used as in Experiment 1. In addition, a new set of 10 color squares was used. For each of the original color squares, there was a new color square from the same hue region as the original but that differed slightly in its brightness and saturation. A listing of the red–green–blue (RGB) values for the 20 colors in the full set of stimuli is provided in Table 5. All other aspects of the stimuli were the same as in Experiment 1.

Procedure—The experimental design and the structure of trials was the same as in Experiment 1, with the following extensions. First, only a single color from each main hue region could appear in the memory set on each trial. The color from each hue region that appeared in the memory set on each trial was chosen at random. On big-change trials, the probe was a randomly chosen color from a new hue region. Only the external condition was tested. In other words, on the big-change trials, the probe was a color from a separate hue region than any of the colors that had appeared in the memory set. On small-change trials,

the probe was a new color from the same hue region as the original. In each change-probability block, half the change trials were big-change trials and half were small-change trials. The ordering of same, small-change and big-change trials was randomized within each block.

The memory-set sizes were 2, 5, and 8, and the change-probability conditions tested on each block were again .3, .5, and .7.

Results

We deleted from the analysis any trial in which the RT was less than 180 ms or greater than 2,500 ms. In addition, for each combination of set size, change probability, stimulus type (same, small-change, big-change), and response type (change vs. no-change), we deleted from the analysis any trial in which the RT was greater than three standard deviations above the mean RT for that combination. These procedures led to deleting at most 2.9% of the trials for any subject.

Before turning to the formal modeling of the individual-subject data, we first provide an overview of the general pattern of results. In particular, in Figure 11 we display the choice-probability and mean RT data averaged across the five subjects. (Symbols show the averaged data, whereas line types show averaged predictions from a to-be-described model.) As can be seen in the top panels, the patterns of choice-probability data on big-change trials and same trials replicate the patterns from Experiment 1: Increases in set size led to lowered accuracy on these trials, whereas increases in change probability led to both increased hit rates and increased false-alarm rates. The effects of set size on small-change trials, however, were smaller than for the big-change and same trials. In addition, subjects were less accurate on the small-change trials than on the big-change trials.

The mean correct RT data are displayed as a function of set size and change probability in the middle panels of the figure. As was the case in Experiment 1, as set size increased, mean correct RTs for big-change and same trials got slower (middle-left panel). The effects of set size on mean correct RTs were smaller for the small-change trials. (Indeed, for small-change trials, average hit RTs were slightly faster for set-size 8 than for set-size 5; however, individual subjects differed in whether they showed the latter pattern.) As shown in the middle-right panel, as change probability increased, hit RTs for both big-change and small-change trials tended to get faster, whereas correct-rejection RTs for same trials got slower. Such results are expected if subjects are more prone to make “change” responses with increasing change-probability. This tendency can also be seen for the mean error RTs shown in the bottom-right panel: As change probability increased, false-alarm RTs for same trials got faster, whereas miss RTs for big-change and small-change trials got slower.

The mean error RT data are shown as a function of set size in the lower-left panel of the figure. For these data, the most important result is that, as predicted by the discrete-slots models, the mean miss RTs on big-change trials are again flat. Note that this pattern was displayed by all five subjects. (The miss RT at set-size-2 is not shown because it is based on an average of only 2.2 observations per subject. Subjects virtually never missed on the big-change trials at the smallest set size.) As described in the introduction to this experiment, the

discrete-slots models no longer predict that false-alarm RTs for the same trials should be a flat function of set size, nor that miss RTs for the small-change trials should be flat. The tendency was for the false-alarm RTs on *same* trials to get faster with increases in set size, and for the miss RTs on *small-change* trials to get slower. (The very slow mean false-alarm RT at set size 2 was due mainly to a single subject.)

Modeling Analyses

The formal models—We fitted straightforward extensions of discrete-slots model 3 (DS3) and continuous-model 2 (C2) to the individual-subject data, using the same methods as already described in Experiment 1. (We focus on these representatives from the discrete-slots and continuous classes because they performed best in Experiment 1. Special cases of the models from each class that did not allow the response thresholds on the memory-based accumulators to vary with set size or change probability again tended to perform worse than did models DS3 or C2.) Applying model DS3 to the present design required the addition of one new free parameter, namely, the mean correct drift rate on the memory-based accumulator on small-change trials. Applying model C2 to the present design required the addition of three new free parameters, namely, the mean correct drift rate on small-change trials at each of the three set-size levels. In all other respects, models DS3 and C2 were the same as applied in Experiment 1.

We also fitted two hybrid models to the data. In hybrid-model H1, we extended discrete-slots model DS3 by assuming that mean correct drift rates on the small-change and same trials varied with set size. In particular, separate small-change and same drift-rate parameters were allowed for each individual set size (the drift-rate parameter on big-change trials, however, was presumed to be invariant with set size, as in Experiment 1.) This hybrid model continues to assume that responding is based on a mixture of cognitive states, namely, memory and guessing. However, the model allows for the possibility that resolution in the memory state changes as memory set-size increases.

The second hybrid model (H2) was a special case of the first and, as explained below, was motivated by the “slots plus averaging” model of Zhang and Luck (2008) and Cowan and Rouders (2009). In particular, model H2 was the same as model H1, except that the mean correct drift rates on small-change and same trials were presumed to be invariant across set sizes 5 and 8. Unique small-change and same drift rates were allowed only for set-size 2. According to the slots-plus-averaging model (e.g., Zhang & Luck, 2008), visual working memory is composed of roughly 3–4 slots, and a set of total resources available to working memory is shared equally among the items that occupy those slots. Note therefore that once the number of to-be-remembered items exceeds the number of slots, then resolution of the items that are stored in the slots will remain fixed as set size increases. However, if the number of to-be-remembered items is *smaller* than the number of discrete slots, then greater resources can be devoted to each individual item. Thus, the slots-plus-averaging model would predict that memory-based drift rates are greater for set-size 2 than for set-sizes 5 and 8 but that the drift rates are invariant across set sizes 5 and 8.

BIC fits—The BIC fits of the models to the five subjects are reported in Table 6. Among the models that presume the operation of mixed states (DS3, H1 and H2), model H2 provides the best BIC fit for all five subjects. In addition, it provides a dramatically better fit than does model C2 for Subjects 1, 2, and 5; essentially the same BIC fit for Subject 3; but a worse BIC fit for Subject 4.

The dramatically better fits of model H2 compared to model C2 for Subjects 1, 2, and 5 provide very clear evidence for the operation of a mixture of memory and guessing states. However, the consistent advantage of model H2 compared to the pure discrete-slots model DS3 provides evidence for a role of continuous changes in resolution as well, at least when memory set size gets very small. Such results are anticipated by the slots-plus-averaging model of Zhang and Luck (2008) and Cowan and Rouders (2009). Although model C2 fares better than does model H2 for Subject 4, we will suggest later that the reasons do not necessarily imply that model C2 is a more appropriate model.

Predictions of averaged data—The predictions from model H2, averaged across all five individual subjects, are displayed along with the observed averaged data in Figure 11. With the exception of the very slow false-alarm RT at the smallest set size (lower left panel), which was due almost entirely to a single subject (S5), model H2 captures all of the summary trends quite well. The explanations for most of its main predictions involving the big-change and same trials are essentially the same as we have already described for the results in Experiment 1. Most important, the model accurately predicts that mean *miss* RTs for the big-change trials are virtually a flat function of set size because those errors arise solely from guessing.

Regarding the small-change trials, the model predicts worse performance on those trials than on the big-change trials (higher error rates and slower RTs) because small-change drift-rate parameters are much smaller than big-change drift-rate parameters. Interestingly, the model also predicts correctly that mean RTs for small-change hits get slightly *faster* at set-size 8. The explanation is as follows. Recall that in the present design, because of the difficulty in discriminating *same* from *small-change* trials, memory-based decisions on such trials may be slow. Indeed, for the present parameter settings, guessing-based hits on small-change trials are faster than are memory-based hits. As set size increases, there is an increased proportion of guessing, so mean hit RTs on the small-change trials get faster.

ROC analyses—As it turned out, model C2 yielded roughly the same predictions of the Figure-8 summary data as did model H2. The reason why model C2 yielded much worse fits to the individual-trials data of some of the individual subjects involved its predictions of their detailed ROC curves.

An example is provided in Figure 12, which shows the observed and predicted ROC curves for Subject 5 (S5). The left panels show the results for the big-change trials and the right panels show the results for the small-change trials. The top panels plot the predictions from model H2, whereas the bottom panels plot the predictions from model C2. Whereas model H2 accounts well for the full set of ROC data, model C2 has some dramatic shortcomings.

In particular, at the largest set size, the predicted ROCs from model C2 fail to come close to the full span of choice probabilities in both the big-change and small-change data.

Why do the models differ in this way? Note that on the big-change trials, in the high change-probability condition, S5's hit rates were nearly unity at every set size. Of course, on trials in which the memory-based accumulation process operated, hit probabilities would be at unity, because of the ease with which the big changes are detected. However, there will also be a substantial proportion of guessing trials that contribute to the hit rates. Thus, because the overall hit probabilities were near unity in the blocks in which change probability was high, S5 apparently adopted an extreme response policy of nearly always guessing "change" if the probed stimulus did not occupy one of the slots. (Note that if responses are based on guessing, then this response policy is the ideal-observer strategy for this design.) Moreover, the subject appears to have also adopted a complementary extreme response policy for guessing "no change" on blocks in which change probability was low.⁶

Taken in combination, these extreme guessing policies lead to ROCs with a large span in the probability space. And although the observed ROCs have slightly more curvature than is predicted by model H2, these large-span ROCs are apparently far closer to linear than continuous model C2 can accommodate, particularly when it is required to simultaneously fit the RT data observed in the task.^{7,8} Subject 1 showed very similar patterns of ROC results, which explains why model C2 fared extremely poorly for that subject as well. Subject 2 (S2) did not show this extreme ROC pattern, and Model C2 can predict S2's ROCs quite well. Thus it is likely that Model H2 is favored for S2 for more subtle reasons involving RTs or the interaction of RTs and accuracy.

Evidence for changes in resolution with set size—As noted earlier, hybrid-model 2 (H2) yielded a better BIC fit than did model DS3 or H1 for all five subjects. This result provides evidence in favor of the idea that, although performance may have involved a mixture of memory-based responding and guessing, resolution within the memory state was different at set-size 2 than at set sizes 5 and 8 (and there was no difference in resolution at the larger set sizes). Indeed, examination of the best-fitting drift-rate parameters from model H2 for each of the individual subjects revealed a highly interpretable pattern. First, as expected, for all subjects the estimated mean correct drift rate on big-change trials was always far greater in magnitude than the drift rates on small-change and same trials. Second, for all subjects, the magnitudes of the small-change and same drift rates at set-size 2 were

⁶Recall that in the present design, false alarm rates on *same* trials would not be near zero even if the subject always responded "no-change" when in the guessing state. False alarms will also be generated from the memory-based process, because of the difficulty of discriminating *same* from *small-change* trials. For the same reason, in the present design, the discrete-slots models can predict slight departures from linear ROCs.

⁷Because memory-based decisions are not made with perfect accuracy in this experiment that includes small-change trials, the DS models no longer predict perfectly linear ROC curves. The degree of predicted departure from linearity will depend on parameter settings from the model.

⁸We conducted numerous additional analyses to investigate the constraints that S5's ROC data imposed on model C2. For example, in one set of analyses, we tested model C2's ability to fit *only* the ROC data, without requiring it to also fit the RTs. However, we imposed the reasonable constraints that (a) the mean correct drift rates in the model would be nonincreasing with set size, that is, memory resolution would not improve as set size increased, and (b) response thresholds would be nondecreasing with increases in set size, that is, the system would not require less evidence for responding at large set sizes compared to small set sizes. Under these constraints, model C2 failed dramatically to fit the ROC data, even in isolation. In another set of analyses, we allowed all of model C2's parameters to vary freely in fitting the ROC data, without imposing the above-stated constraints. Although it could then fit the ROC data, with those parameter settings, it made bizarre predictions of the patterns of RT data.

always greater than or equal to those at the larger set sizes, indicating greater memory resolution at set-size 2. This pattern of results is precisely as predicted by the slots-plus-averaging model (Cowan & Rouder, 2009; Zhang & Luck, 2008).

Continuous subjects—As we noted earlier, model C2 yielded a better BIC fit than did model H2 for Subject 4 (and the models yielded roughly equivalent BIC fits for Subject 3). However, our detailed examinations indicated that both models appeared to yield very good accounts of the mean correct and error RT data as well as the ROC data of these individual subjects. The main reason why Subject 3 (S3) and Subject 4 (S4) differed from Subject 1 (S1) and S5 is that S3 and S4 produced very small changes in hit and false-alarm rates as a function of the manipulations in change probability across blocks. Thus, the three points that defined each ROC curve tended to lie close together (especially those of S4), making it difficult to discriminate between the predictions and the models. The BIC advantage for model C2 for these subjects thus arises because it uses fewer free parameters, not because it provides a better absolute fit to their data.⁹ In our view, the results from these subjects are relatively nondiagnostic with respect to choosing between members of the continuous and discrete-slots families, rather than posing major challenges to the discrete-slots family. A simplicity bias in model selection is an excellent policy and should apply if considering S4 in isolation. However, model H2 provides a good and coherent account of the data from the full set of subjects, so that model is strongly favored when considering the group as a whole (cf. Jang, Wixted, & Huber, 2011).

Summary

Overall, models that make allowance for the mixture of states (memory plus guessing) assumed by the discrete-slots models provide better accounts of the Experiment-2 data than does a model based solely on continuous sharing of resources among all items of the memory set. Thus, the requirement that subjects make fine-grained discriminations between study items and similar lures does not seem to have led to a major change in processing modes across the two experiments. However, the present design also provided evidence for a hybrid account of performance in the task, in which resolution for items stored in the discrete slots improved when memory set size grew very small. Such results are anticipated by the “slots-plus-averaging” models of Zhang and Luck (2008) and Cowan and Rouder (2009).

General Discussion

Summary

In sum, in this research we developed two families of models to account for both accuracies and RTs in tasks of visual WM change detection. The main variables manipulated in the tasks included memory set size, change probability, and whether changes were big or small in magnitude. One family of models formalized the discrete-slots view, which holds that an object either does or does not occupy a slot in visual WM. Objects that occupy one of the

⁹Using the alternative Akaike information criterion (AIC) statistic as a criterion of fit, which places a smaller penalty on number of free parameters than does the BIC statistic, model H2 yields a much better fit than does model C2 for all five subjects. Indeed, even model DS3 produces a better AIC fit than does model C2 for all five subjects, although its advantage for Subject 4 is minuscule.

slots are stored with constant resolution. For objects that do not occupy one of the slots, there is a complete loss of resolution, and observers are forced to guess. The second family formalized the continuous shared-resources view, which holds that as memory set size increases, the observer stores lower resolution memories of all of the members from the memory set. Hybrid models that combined assumptions from the two families were also examined.

The key new idea from the discrete-slots family is that RTs are presumed to arise from a mixture of basis distributions, one that reflects memory-based evidence accumulation, and a second that reflects guessing-based accumulation. Different versions of the discrete-slots models varied the detailed assumptions by which the two types of accumulation processes operated across conditions.

Under the present experimental conditions, the data pointed decidedly in favor of the discrete-slots view, although with the refinement at very small set sizes, there is increased resolution for items in the memory state. In general, the continuous shared-resource models had difficulty accounting jointly for the correct and error RTs (Experiment 1) or else failed to account for the detailed structure of ROC curves in cases in which subjects appeared to adopt extreme guessing policies (Experiment 2). The discrete-slots models did a better job by producing better quantitative accounts of the choice probability and RT data, including detailed accounts of the distributions of RTs, and by also matching certain key qualitative predictions. For example, the discrete-slots models predicted correctly that certain types of error RTs would be nearly invariant with set size. It remains an open question whether evidence favoring discrete-slots processing will continue to be observed under alternative experimental conditions. For example, the conclusions we have reached are of course limited to tasks using masks. Our tasks and others investigating the present issues use form and color masking to eliminate very short-term visual memory representations such as the iconic memory first studied by Sperling (1960). It remains an open question what the results would be if alternative masking procedures were used.

Inattention, Fast Guessing, and Error RTs

Although the discrete-slots models performed well overall, perhaps their main limitation is that they failed to account for the fast false-alarm RTs observed at the smallest set size in Experiment 1. Although those results were based on very small sample sizes (because observers almost never erred at the smallest set size), the pattern was consistent across almost all of the subjects.

We think it is likely that these low-frequency events reflect additional processes that go beyond all of the models as they are presently formalized. We should note that even when restricted to predicting choice-probability data, Rouder et al. (2008) argued for the importance of extending certain versions of discrete-slots models by incorporating an “inattention” parameter (see also Rouder, Morey, Morey, & Cowan, 2011). In these extended models, the observer is presumed to attend to the stimulus display with probability a and thus fails to attend with probability $1-a$. On those (rare) trials in which the observer fails to attend, she is forced to guess. Following Rouder et al. (2008), a straightforward extension of the present discrete-slots RT models would be to include an analogous

inattention process, while making the further assumption that the time-course of processing on such trials differs from those on attended trials. In a nutshell, we could posit a fast “guess-change” process, analogous to hypotheses involving fast-guess processes that have been considered by other researchers (e.g., Link, 1982; Ollman, 1966; Yellott, 1971). Importantly, because the fast-guess process occurs rarely, it would have a minuscule influence on the pattern of RTs associated with the larger set sizes, where the standard guess-accumulation process (right branch of Figure 2) would dominate guessing behavior. However, on trials involving the very small set sizes, observers virtually always enter the memory state, so the standard guess-accumulation process rarely operates. On these trials, the special fast-guess process would contribute a substantial proportion of the total guessing observations, thereby having a big influence on the pattern of error RTs.

We conducted exploratory analyses of such extended fast-guess models for the present data. Because so few total observations were involved, it is perhaps not surprising that these extensions did not improve the models’ BIC fits to the data. However, we have verified that including such a process allows the models to capture the very fast mean false-alarm RTs at the smallest set size, while leaving predictions for all of the other data points virtually unchanged.

Illusory Conjunctions, Feature Migrations, and Error RTs

The possibility of inattention and fast-guessing does not explain, however, why faster-than-expected RTs at the smallest set size tended to occur for false alarms but not misses (in Experiment 1). An explanation for this more nuanced result may reside in the extensive literature that deals with the phenomenon of “illusory conjunctions” (e.g., Prinzmetal, Diedrichsen, & Ivry, 2001; Treisman & Schmidt, 1982; Wolford & Shum, 1980).

The basic idea is that, on some small proportion of trials, a feature from one object or location may migrate to another object or location in the visual display. Suppose that this type of feature migration took place in our design. The process would cause errors not because the subject is guessing due to an empty slot but because the feature stored in the slot is incorrect (cf. Bays et al., 2009). Furthermore, because the memory-based accumulation process operates faster than the guessing process when changes are big, those types of errors would tend to be fast. Because there are hardly any true guesses at the small set sizes, the feature-migration trials would contribute a large proportion of the errors in those conditions, thereby explaining the fast false-alarm RTs. Moreover, as we explain in Appendix A, such a process would also predict an asymmetry in results for false alarms versus misses because misses would tend not to arise from feature migrations. We present preliminary empirical evidence in Appendix A that such a feature-migration process may indeed be part of a complete explanation of performance in these tasks. Again, however, we decided not to append such a process to the formalized models in the present article because it would have added a great deal of complexity to handle data with very few observations.

Extensions to Continuous Shared-Resources Models

In very recent work, van den Berg et al. (2012) introduced a new type of continuous shared-resource model of visual WM that captures what appears to be “guessing” in human

responses. The key idea is that continuous resources are not only variable across trials but are also variable across the items within a given memory set. If an observer devotes minimal resources to one of the memory-set items and that item is probed, it will appear as if the subject is guessing.

There are a wide variety of ways in which one might formalize variable resource allocation across the items of a memory set in the present LBA-based RT models. Borrowing from findings of Donkin and Nosofsky (2012a), we implemented and tested one particular version of such a model. In particular, Donkin and Nosofsky (2012a) conducted a short-term probe-recognition experiment involving the sequential presentation of memory-set items. In modeling the choice and RT data from that experiment, Donkin and Nosofsky (2012a) obtained evidence that the “memory strength” of the items decreased as a power function of their lag of presentation (see also Anderson & Schooler, 1991; Wickelgren, 1974; Wixted & Ebbesen, 1991). In the context of the present types of visual WM tasks, one might imagine an analogous process in which the observer shifts covert visual attention across the memory-set items, with more recently attended items having greater memory strength. To approximate the idea, we assumed that mean correct drift rates associated with the items in a set varied according to a power function of their “covert lag” (see Appendix B for details). Thus, analogous to the ideas from van den Berg et al. (2012), the drift rates in the change and no-changes accumulators were assumed to be doubly stochastic. There was variability in mean drift rates across the items within a given set (that followed a power-function distribution) but also momentary variability around those mean drift rates as in standard LBA modeling (captured by the standard-deviation-of-drift-rate parameter σ).

Note that this variable-resource model is a mixture model. For any given set size, the predicted RT arises from a probabilistic mixture of component RTs associated with each of the mean drift rates. Also note that with appropriate choice of the power-function free parameters, the magnitude of many of the mean drift rates can be set such that mean drift rate on the correct accumulator is virtually identical to mean drift rate on the incorrect accumulator, that is, what is essentially guessing behavior. Thus, just as is the case for the discrete-slots models, this variable-resources model both is a mixture model and can incorporate guessing. Despite this flexibility, we found that this variable-resources model continued to provide substantially worse fits to our RT-distribution data than did the two best-fitting discrete-slots models (DS2 and DS3; see Appendix B for details).

Perhaps the main conceptual difference between the discrete-slots models and this variable-resources model is that memory-based responding and “guessing-based” responding are governed by *separate* sets of accumulators in the former, but by a common set of accumulators in the latter. Thus, the parameters that influence the evidence-accumulation process, such as start-point variability and response-threshold settings, must be common for all items in the variable-resources model. By contrast, separate sets of parameters are presumed to govern memory-based accumulation and guessing-based accumulation according to the discrete-slots models. These separate forms of evidence accumulation are apparently reflected in the detailed RT-distribution data that served as targets for the present modeling. These modeling results may reflect that although accuracies associated with guessing can be accounted for by assuming that an item has received minimal resource

allocation, the time course of guessing follows a different path than does the time course of memory-based decision making.

Relations to Other Mixture Models

The present mixed-state (discrete-slots) models share certain properties with mixture models for RT distributions developed by Ratcliff and Tuerlinckx (2002). (Their mixture models were not intended to account specifically for performance in visual WM change-detection tasks, but rather as a general approach to modeling performance in a wide variety of experimental paradigms.) In an attempt to account for “contaminant” RTs, these researchers suggested that on some small proportion of trials, a response and RT may be generated from a process that is separate from the main evidence-accumulation process that governs behavior. For simplicity, Ratcliff and Tuerlinckx (2002) assumed that this contaminant process generates RTs that are uniformly distributed. There are several differences between the mixed-state models developed here and the contaminant approach from Ratcliff and Tuerlinckx. First, the present modeling efforts build in a structure for how the probability of entering into the guess state is expected to vary across experimental conditions. By contrast, in Ratcliff and Tuerlinckx’s modeling, there is no structure for how the probability of the contaminant process is expected to vary across conditions. Second, in the present approach, the guessing process is itself modeled in terms of an accumulation process, rather than simply being specified as obeying some descriptive distribution. Third, the parameters of that accumulation process are posited to vary systematically across experimental conditions. Thus, the forms of the RT distributions that are generated from guessing vary systematically across conditions as well. In short, the present models attempt to characterize the “cognitive psychology” of guessing in a manner that goes well beyond the contaminant mechanism developed by Ratcliff and Tuerlinckx.

Toward More Parsimonious Discrete-Slots Models

Our central goal in this initial theoretical inquiry was to develop and contrast general versions of the discrete-slots and continuous shared-resources models of visual working-memory RTs. To achieve this goal of generality, we decided not to impose a variety of parameter constraints that could lead to the development of more parsimonious versions of each model class. In this section we outline routes of future research aimed at achieving this latter goal.

First, in the discrete-slots models tested in this article, we allowed the memory-state probabilities m_i to vary as free parameters. An alternative approach is to constrain those parameters by imposing assumptions involving fixed memory capacity across the different set-size conditions. For example, in fitting their discrete-slots models to ROC data, Rouder et al. (2008) denoted the number of available slots by K and assumed that the value of K was constant across the different memory set sizes N . Thus, across the different set size conditions, they assumed that the probability that any given study item would occupy one of the slots when memory was probed would be given by $m = \min(K/N, 1)$. Because they estimated fractional values of K in fitting the model, they were clearly assuming that the number of available slots in visual WM was variable across trials but that the *mean* of this variable number of slots was fixed across set sizes. (Also, because they assumed that the

memory-state probability m was equal to one when $K > N$, their implicit assumption was that the lower limit of the variable distribution must be at least equal to N .) This approach to modeling visual WM would allow a reduction in the number of free parameters that we used for fitting the present discrete-slots models to our RT data. However, as we described earlier in our General Discussion, Rouder et al. (2008, 2011) have emphasized an apparent role of attention processes as well, which further modulate the probabilities that items will reside in the memory state. In short, in cases involving only three different memory set sizes, there is not much difference in number of free parameters required by our “free m_i ” approach versus the fixed-capacity approach applied by Rouder et al. (2008). A greater savings in relative number of memory-state parameters would be achieved, however, in designs that tested a larger number of set sizes.

Another potential route to developing more parsimonious versions of the discrete-slots models is to take advantage of our findings that the parameter estimates for almost all of our subjects varied in highly systematic and psychologically meaningful ways across conditions. For example, as change probability increased, participants set more lenient thresholds on the guess-change accumulator and stricter thresholds on the guess-no-change accumulator. Thus, rather than estimating each threshold as a free parameter, future research might impose functional constraints on how the magnitudes of the thresholds vary with change probability. Still another approach to achieving greater parsimony is to develop deeper process-level accounts of how the drift rates in the memory-based accumulators may vary with manipulations of change magnitude (e.g., Fific, Little, & Nosofsky, 2010).¹⁰

Unpacking the Gating Process

To formalize the discrete-slots models in the present investigation, we assumed an initial gating process that informed the system whether memory-based information was available in the probed location of the visual display. We represented the probabilistic outcome of the gating process by the free parameters m_i but did not provide a mechanistic account of the process. In our view, it would be straightforward to append such a mechanism to the discrete-slots models, but this approach would have little bearing on the evaluation of the competing models under the current experimental conditions. For example, it is easy to imagine an initial LBA mechanism at the “gate” in Figure 2, which is responsible for determining only whether memory-based information is or is not present at the probed location of the visual display. Such a mechanism would not be responsible, however, for determining whether the probed location changed from study to test. This type of distinction dates back at least to Garner (1974), who distinguished between *state* limitations versus *process* limitations in information processing. State limitations are concerned only with

¹⁰In addition, in recent work, Hyun, Woodman, Vogel, Hollingworth, and Luck (2009) conducted initial empirical investigations of the time course of processing in certain types of visual change-detection tasks. In particular, they explored how the time course of manual RTs, saccades, and event-related potential components varied with specific task goals and set-size manipulations. There were many important differences between the Hyun et al. tasks and the tasks reported in this article. For example, because Hyun et al.’s goals did not involve testing for evidence of mixed-state processing, the set sizes tested in their research were always less than or equal to four. In addition, rather than testing memory with a probe in only a single location, their test displays presented probes in all of the locations that had appeared in the study displays, thereby requiring observers to engage in simultaneous memory and visual search processes to perform the change-detection tasks. Hyun et al. found evidence for automatic, unlimited capacity processing for some aspects of their change-detection tasks, but for limited-capacity processing in other aspects. Their findings could help constrain a fully developed theory of the memory-accumulation processes that operate in the discrete-slots models developed here.

whether information is present or absent. By contrast, in process-based limitations, information is present, but the system is limited in its ability to discriminate which of two or more events has occurred. Future research might be aimed at conducting manipulations to unpack the mechanisms that underlie the “presence-absence” gating process. For example, one might vary the timing of the cue that signals the observer regarding the location of the probe, or one might vary the intensity of the probe itself.

Another interesting avenue for future research might involve making direct queries to the subject about the outcome of the gating process. If one makes the strong assumption that observers have full conscious access to the outcome of the gating process, then they should be able to report on each trial whether their decisions were based on memory or guessing. This approach would allow one to make extremely strong predictions regarding the substructure of the data.¹¹ For example, on trials in which subjects report they were guessing, then not only should error-RT distributions be invariant with set size, but correct-RT distributions should be invariant with set size as well. It is a wide open question, however, the extent to which observers do indeed have full conscious access to the outcome of the gating process.

Extensions to Other Perceptual and Cognitive Tasks

The focus of the present work was on visual WM change detection, a domain where discrete-slots models that posit a mixture of information-based and guessing-based cognitive states have played a prominent role. The present mixed-state models of choice and RT, however, have a potentially far wider range of applicability. There are numerous types of perceptual and cognitive tasks for which it seems plausible that a mixture of information-based and guessing-based cognitive states might operate. Applications of the present types of mixed-state RT models could yield greater insights into the cognitive processes that operate in those domains as well.

For example, consider a task of perceptual identification, in which a single perceptual object is presented under degraded conditions, and an observer is required to rapidly identify or classify it. Common techniques for producing error-prone identification include using short display durations and/or masking. It seems plausible that on some trials, the observer may be able to extract at least partial information from the degraded object and use that information as a basis for the classification decision. But on other trials, there may be a failure to reach threshold, and the system is forced to guess. Application of the present mixed-state models of choice and RT may yield diagnostic information in support of this hypothesis.

As a second example, consider the domain of long-term recognition memory. A classic idea is that two distinct processes—familiarity and recollection—contribute to recognition judgments. Furthermore, according to certain versions of these dual-process models, familiarity-based and recollection-based paths to recognition follow distinct mental routes. Many past efforts at evaluating alternative types of single-process and dual-process models have relied on detailed modeling of ROC curves (e.g., Dube & Rotello, 2012; Wixted, 2007; Yonelinas, 1994). However, in recent work, researchers have developed and tested single-

¹¹We thank John Wixted for suggesting to us this insightful idea.

process models on their ability to account jointly for choice probability, confidence, and *RT* data in long-term recognition tasks (e.g., Dube, Starns, Rotello, & Ratcliff, 2012; Ratcliff & Starns, 2009; Rotello & Zeng, 2008; Starns, Ratcliff, & McKoon, 2012). A theme that has emerged from this recent literature is that ROC data and RT data should be considered jointly when interpreting the underlying memory processes that are involved. Another important approach that would complement these recent efforts may be to formalize dual-route models of choice and RT in such tasks. For example, in such models, one process might involve accumulation of familiarity-based information, whereas a second would involve accumulation of recollection-based information. The question is whether this alternative type of mixed state, dual-route model might yield a better account of the complete set of recognition choice probabilities and RTs than do members from the other classes. If so, then such a result would provide an interesting form of evidence in favor of the dual-route approaches.

Acknowledgments

This work was supported by an Australian Research Council's Discovery Projects funding scheme (DP120100124) and an Australian Research Council Early Career Award (DP130100124) to Christopher Donkin, Grant R01-EY019265 from the National Institutes of Health to Jason M. Gold, and Grant FA9550-12-1-0255 from the Air Force Office of Scientific Research to Richard M. Shiffrin.

References

- Alvarez GA, Cavanagh P. The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*. 2004; 15:106–111.10.1111/j.0963-7214.2004.01502006.x [PubMed: 14738517]
- Anderson JR, Schooler LJ. Reflections of the environment in memory. *Psychological Science*. 1991; 2:396–408.10.1111/j.1467-9280.1991.tb00174.x
- Awh E, Barton B, Vogel EK. Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science*. 2007; 18:622–628.10.1111/j.1467-9280.2007.01949.x [PubMed: 17614871]
- Barton B, Ester EF, Awh E. Discrete resource allocation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*. 2009; 35:1359–1367.10.1037/a0015792 [PubMed: 19803642]
- Bays PM, Catalao RFG, Husain M. The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*. 2009; 9:1–11.10.1167/9.10.7 [PubMed: 19810788]
- Bays PM, Gorgoraptis N, Wee N, Marshall L, Husain M. Temporal dynamics of encoding, storage, and reallocation of visual working memory. *Journal of Vision*. 2011; 11:1–15.10.1167/11.10.6
- Bays PM, Husain M. Dynamic shifts of limited working memory resources in human vision. *Science*. 2008; 321:851–854.10.1126/science.1158023 [PubMed: 18687968]
- Brown SD, Heathcote A. The simplest complete model of choice reaction time: Linear ballistic accumulation. *Cognitive Psychology*. 2008; 57:153–178.10.1016/j.cogpsych.2007.12.002 [PubMed: 18243170]
- Cowan N. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*. 2001; 24:87–114.10.1017/S0140525X01003922 [PubMed: 11515286]
- Cowan N, Elliott EM, Saults JS, Morey CC, Mattox S, Hism-jatullina A, Conway ARA. On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*. 2005; 51:42–100.10.1016/j.cogpsych.2004.12.001 [PubMed: 16039935]
- Cowan N, Roudier JN. Comment on “Dynamic shifts of limited working memory resources in human vision”. *Science*. 2009; 323:877.10.1126/science.1166478 [PubMed: 19213899]

- Donkin C, Brown SD, Heathcote A, Wagenmakers EJ. Diffusion versus linear ballistic accumulation: Different models but the same conclusions about psychological processes? *Psychonomic Bulletin & Review*. 2011; 18:61–69.10.3758/s13423-010-0022-4 [PubMed: 21327360]
- Donkin C, Nosofsky RM. A power-law model of psychological memory strength in short- and long-term recognition. *Psychological Science*. 2012a; 23:625–634.10.1177/0956797611430961 [PubMed: 22527527]
- Donkin C, Nosofsky RM. The structure of short-term memory scanning: An investigation using response-time distribution models. *Psychonomic Bulletin & Review*. 2012b; 19:363–394.10.3758/s13423-012-0236-8 [PubMed: 22441957]
- Dube C, Rotello CM. Binary ROCs in perception and recognition memory are curved. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2012; 38:130–151.10.1037/a0024957
- Dube C, Starns JJ, Rotello CM, Ratcliff R. Beyond ROC curvature: Strength effects and response time data support continuous-evidence models of recognition memory. *Journal of Memory and Language*. 2012; 67:389–406. [PubMed: 22988336]
- Fific M, Little DR, Nosofsky RM. Logical-rule models of classification response times: A synthesis of mental-architecture, random-walk, and decision-bound approaches. *Psychological Review*. 2010; 117:309–348.10.1037/a0018526 [PubMed: 20438229]
- Garner, WR. The processing of information and structure. Potomac, MD: Erlbaum; 1974.
- Hyun JS, Woodman GF, Vogel EK, Hollingworth A, Luck SJ. The comparison of visual working memory representations with perceptual inputs. *Journal of Experimental Psychology: Human Perception and Performance*. 2009; 35:1140–1160.10.1037/a0015019 [PubMed: 19653755]
- Jang Y, Wixted JT, Huber DE. The diagnosticity of individual data for model selection: Comparing signal-detection models of recognition memory. *Psychonomic Bulletin & Review*. 2011; 18:751–757.10.3758/s13423-011-0096-7 [PubMed: 21538201]
- Link SW. Correcting response measures for guessing and partial information. *Psychological Bulletin*. 1982; 92:469–486.10.1037/0033-2909.92.2.469
- Loftus GR, Masson MEJ. Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*. 1994; 1:476–490.10.3758/BF03210951 [PubMed: 24203555]
- Luck SJ, Vogel EK. The capacity of visual working memory for features and conjunctions. *Nature*. 1997; 390:279–281.10.1038/36846 [PubMed: 9384378]
- Macmillan, NA.; Creelman, CD. Detection theory: A users' guide. Cambridge, England: Cambridge University Press; 1991.
- McElree B, Doshier BA. Serial position and set size in short term memory: The time course of recognition. *Journal of Experimental Psychology: General*. 1989; 118:346–373.10.1037/0096-3445.118.4.346
- Meyer DE, Yantis S, Osman AM, Smith JEK. Temporal properties of human information processing: Tests of discrete versus continuous models. *Cognitive Psychology*. 1985; 17:445–518.10.1016/0010-0285(85)90016-7 [PubMed: 4064623]
- Nosofsky RM, Little DR, Donkin C, Fific M. Short-term memory scanning viewed as exemplar-based categorization. *Psychological Review*. 2011; 118:280–315.10.1037/a0022494 [PubMed: 21355662]
- Ollman RT. Fast guesses in choice reaction time. *Psychonomic Science*. 1966; 6:155–156.
- Pashler H. Familiarity and visual change detection. *Perception & Psychophysics*. 1988; 44:369–378.10.3758/BF03210419 [PubMed: 3226885]
- Prinzmetal W, Diedrichsen J, Ivry RB. Illusory conjunctions are alive and well: A reply to Donk (1999). *Journal of Experimental Psychology: Human Perception and Performance*. 2001; 27:538–541.10.1037/0096-1523.27.3.538 [PubMed: 11424643]
- Ratcliff R. A theory of memory retrieval. *Psychological Review*. 1978; 85:59–108.10.1037/0033-295X.85.2.59
- Ratcliff R, Starns JJ. Modeling confidence and response time in recognition memory. *Psychological Review*. 2009; 116:59–83.10.1037/a0014086 [PubMed: 19159148]

- Ratcliff R, Tuerlinckx F. Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*. 2002; 9:438–481.10.3758/BF03196302 [PubMed: 12412886]
- Ratcliff R, Van Zandt T, McKoon G. Connectionist and diffusion models of reaction time. *Psychological Review*. 1999; 106:261–300.10.1037/0033-295X.106.2.261 [PubMed: 10378014]
- Rotello CM, Zeng M. Analysis of RT distributions in the remember-know paradigm. *Psychonomic Bulletin & Review*. 2008; 15:825–832.10.3758/PBR.15.4.825 [PubMed: 18792511]
- Rouder JN, Morey RD, Cowan N, Zwilling CE, Morey CC, Pratte MS. An assessment of fixed-capacity models of visual working memory. *Proceedings of the National Academy of Sciences of the United States of America*. 2008; 105:5975–5979.10.1073/pnas.0711295105 [PubMed: 18420818]
- Rouder JN, Morey RD, Morey CC, Cowan N. How to measure working memory capacity in the change detection paradigm. *Psychonomic Bulletin & Review*. 2011; 18:324–330.10.3758/s13423-011-0055-3 [PubMed: 21331668]
- Schwarz G. Estimating the dimension of a model. *The Annals of Statistics*. 1978; 6:461–464.10.1214/aos/1176344136
- Sperling G. The information available in brief visual presentations. *Psychological Monographs, General and Applied*. 1960; 74(11 Serial No 498):1–29.
- Starns JJ, Ratcliff R, McKoon G. Evaluating the unequal-variability and dual-process explanations of zROC slopes with response time data and the diffusion model. *Cognitive Psychology*. 2012; 64:1–34.10.1016/j.cogpsych.2011.10.002 [PubMed: 22079870]
- Treisman A, Schmidt H. Illusory conjunctions in the perception of objects. *Cognitive Psychology*. 1982; 14:107–141.10.1016/0010-0285(82)90006-8 [PubMed: 7053925]
- van den Berg R, Shin H, Chou WC, George R, Ma WJ. Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences of the United States of America*. 2012; 109:8780–8785.10.1073/pnas.1117465109 [PubMed: 22582168]
- Vogel EK, Woodman GF, Luck SJ. Storage of features, conjunctions, and objects in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*. 2001; 27:92–114.10.1037/0096-1523.27.1.92 [PubMed: 11248943]
- Wickelgren WA. Single-trace fragility theory of memory dynamics. *Memory & Cognition*. 1974; 2:775–780.10.3758/BF03198154 [PubMed: 24203753]
- Wilken P, Ma WJ. A detection theory account of change detection. *Journal of Vision*. 2004; 4:1120–1135.10.1167/4.12.11 [PubMed: 15669916]
- Wixted JT. Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*. 2007; 114:152–176.10.1037/0033-295X.114.1.152 [PubMed: 17227185]
- Wixted JT, Ebbesen EB. On the form of forgetting. *Psychological Science*. 1991; 2:409–415.10.1111/j.1467-9280.1991.tb00175.x
- Wolford G, Shum KH. Evidence for feature perturbations. *Perception & Psychophysics*. 1980; 27:409–420.10.3758/BF03204459 [PubMed: 7383829]
- Yellott JJ. Correction for fast guessing and the speed-accuracy tradeoff in choice reaction time. *Journal of Mathematical Psychology*. 1971; 8:159–199.10.1016/0022-2496(71)90011-3
- Yonelinas AP. Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1994; 20:1341–1354.10.1037/0278-7393.20.6.1341
- Zhang W, Luck SJ. Discrete fixed-resolution representations in visual working memory. *Nature*. 2008; 453:233–235.10.1038/nature06860 [PubMed: 18385672]

Appendix A. Feature Migrations: Fast False Alarms Versus Fast Misses

Why would the feature-migration process lead to an asymmetry in results for false alarms versus misses at the smallest set size? First, consider the external condition at set-size 2. (Recall that in the external condition, “change” probes were always colors from outside the study display.) If a location is probed with a no-change stimulus, but a feature-migration

occurred, it will result in a fast false-alarm RT (i.e., a “change” response) for the reason described in the text. But if a location is probed with a change stimulus, there will never be a case in which a miss (a “no-change” response) will occur because of a feature migration. If a feature migration has occurred, the observer will still respond “change,” because the probe is different from any of the study colors. In sum, in the external condition, feature migrations will cause fast false alarms but will not cause fast misses.

The argument is more involved for the internal condition, in which “change” probes are always colors from inside the study display. In our internal condition, the smallest set size was equal to 3. Suppose that a feature migration occurs at the probed location. If the probe is a no-change stimulus, then as already explained, that feature migration would cause a fast false-alarm RT. However, if the probe is a change stimulus, it could now cause a fast miss RT. For example, suppose the probed study location was red but the feature green migrated to that location. Suppose further that the test probe was green. In this case, there would be a fast “no-change” response (i.e., a miss). Even so, averaged across trials, the effect would tend to be greater for false alarms than for misses. The reason is that only a single type of feature migration can result in a miss (i.e., one in which the migrating feature happens to be identical to the probe). By contrast, feature migrations from any of the locations in the display will result in false alarms. Thus, as long as set size is 3 or greater, there are more opportunities for feature migrations to result in false alarms than misses.

Prior to conducting Experiment 1, we tested a group of 96 subjects in a pilot study. The design was similar to Experiment 1, with the main difference that each subject participated for only a single session. Half participated in an external-change condition and half in an internal-change condition. Averaged across subjects, the mean RT data for false alarms and misses were very similar to the pattern predicted above. Mean false-alarm and miss RTs were nearly equal across set-size-5 and set-size-8 in both the external and internal conditions. False-alarm RTs were faster for set-size-3 compared to the larger set sizes in both the external and internal conditions. Interestingly, however, whereas miss RTs for set-size-3 were somewhat faster than for the larger set sizes in the internal condition, the miss RTs were nearly flat across all set sizes in the external condition. However, this last fine-grained interaction effect did not reach statistical significance.

In sum, the patterns of faster-than-expected error RTs observed at the very small set sizes appear to be generally consistent with the predictions from the discrete-slots models, assuming that there is some small probability that feature migrations take place.

Appendix B. Variable Shared-Resources Model

The variable shared-resources model (VSRM) defines a distribution of mean correct drift rates, with a separate mean correct drift rate associated with each item of a memory set with size m . In particular, the mean correct “change” drift rate associated with an item that has “covert lag” k in a memory set of size m is given by

$$v_C(m, k) = .5 + \alpha(C, m)k^{-\beta(C, m)}, \quad (\text{B1})$$

where $\alpha(C,m)$ [$m = (2 \text{ or } 3), 5, 8$] and $\beta(C,m)$ [$m = (2 \text{ or } 3), 5, 8$] are free parameters that define the power function that relates “change” drift rate to set size and covert lag. Each item in the memory set is presumed to have a unique covert lag k , where k ranges from 1 to m . In the present application, the constant .5 was added to the power-function computation so that the mean correct drift rate would always be greater than or equal to the mean incorrect drift rate. (As is the case in standard LBA modeling, the mean incorrect drift rate is given by $v_C' = 1 - v_C$). Analogously, the mean correct “no change” drift rate on the no-change accumulator is given by

$$\nu_{NC}(m,k) = .5 + \alpha(NC,m)k^{-\beta(NC,m)}, \quad (B2)$$

where $\alpha(NC,m)$ and $\beta(NC,m)$ are the free parameters that define the no-change power functions. Besides defining a variable distribution of mean drift rates, the important conceptual point is that as covert-lag k grows larger, the mean correct and incorrect drift rates converge toward .5, so there is no differential memory-based information for making change versus no-change responses. Analogous to the ideas advanced by van den Berg et al. (2012), this mechanism can lead to apparent “guessing.”

In all other respects, the VSRM is the same as continuous shared-resources model 2 (C2), defined in the main text. Note that in the special case in which the β parameters are set equal to zero for all set sizes m , there is no reduction in drift rate with covert lag, and the VSRM reduces to model C2. Allowing the α and β parameters to vary freely, the VSRM uses an additional six free parameters compared to model C2.

The BIC scores yielded by fitting the VSRM to the data were (for Subjects 1–8, respectively): $-5,497.5$, $-4,090.6$, $-4,187.0$, $-3,216.9$, $1,617.7$, $-7,471.9$, $-1,451.3$, and $-1,213.4$. The reader may verify that, with the exception of Subject 2, the BIC fit for the VSRM was always worse than for model C2 (compare to Table 2). (In addition, the VSRM yielded a worse fit to S2’s data than did discrete-slots models D2 and D3). Thus, this approach to formalizing variable shared resources within the framework of the continuous models was not successful in accounting for our data. Undoubtedly, there are other approaches to formalizing continuous models of visual WM response times that make allowance for variable shared resources. The fits yielded by the discrete-slots models should provide a challenging yardstick to gauge such alternative approaches.

Figure 1.

A: Illustration of the procedure in the visual WM change-detection task. B: The form of the isosensitivity and isobias curves predicted by the Equation-1 discrete-slots model. From “An Assessment of Fixed-Capacity Models of Visual Working Memory,” by J. N. Rouder, R. D. Morey, N. Cowan, C. E. Zwillig, C. C. Morey, and M. S. Pratte, 2008, *Proceedings of the National Academy of Sciences of the United States of America*, 105, p. 5976. Copyright 2008 by National Academy of Sciences. WM = working memory.

Figure 2.

Schematic illustration of the mixed memory-based and guessing-based accumulation processes that underlie the discrete-slots models of visual WM choice and response time. WM = working memory; m_i = memory-state probability at set-size i .

Figure 3.

Schematic illustration of the LBA model of the memory-based accumulation process for the change-detection task. LBA = linear ballistic accumulator; R_c = response-threshold parameter on memory-based change accumulator; R_{NC} = response-threshold parameter on memory-based no-change accumulator; V_{NC} = mean correct drift rate on memory-based no-change accumulator; V_c = mean correct drift rate on memory-based change accumulator.

Figure 4.

Experiment 1: Averaged observed and predicted choice-probability and mean response time (RT) data. Observed data are shown by different symbol types. Predictions from the discrete-slots-3 (DS3) model are shown by different line types. The error bars that are illustrated are within-subjects error bars following the methods of Loftus and Masson (1994). (To reduce clutter, error bars are not shown for the correct RT data.) RTs are measured in seconds.

Figure 5.

Choice probability and mean response time (RT) data for Subject 3 (S3) of Experiment 1 plotted as a function of set size and change probability. Predictions from discrete-slots model 2 (DS2).

Figure 6.

Choice probability and mean response time (RT) data for Subject 3 (S3) of Experiment 1 plotted as a function of set size and change probability. Predictions from continuous model 2 (C2).

Figure 7.

Choice probability and mean response time (RT) data for Subject 1 (S1) of Experiment 1 plotted as a function of set size and change probability. Predictions from discrete-slots model 3 (DS3).

Figure 8.

Choice probability and mean response time (RT) data for Subject 1 (S1) of Experiment 1 plotted as a function of set size and change probability. Predictions from continuous model 2 (C2).

Figure 9.

Experiment 1: Observed (histograms) and predicted (smooth curves) response time (RT) distribution data for each combination of change probability, set size, and response type (hits, misses, false alarms, correct rejections). Results for Subject 1 (S1) and discrete-slots model 3 (DS3). Prob = probability; Cor. = correct.

Figure 10.

Experiment 1: Observed (histograms) and predicted (smooth curves) response time (RT) distribution data for each combination of change probability, set size, and response type (hits, misses, false alarms, correct rejections). Results for Subject 3 (S3) and discrete-slots model 2 (DS2). Prob = probability; Cor. = correct.

Figure 11.

Experiment 2: Averaged observed and predicted choice-probability and mean response time (RT) data. Observed data are shown by different symbol types. Predictions from hybrid-model 2 (H2) are shown by different line types. RTs are measured in seconds. P = probability.

Figure 12.

Observed and predicted receiver operating characteristic curves for Subject 5 of Experiment 2. Top: Predictions from hybrid-model 2 (H2). Bottom: Predictions from continuous-model 2 (C2). Left: Big-change trials. Right: Small-change trials. Open symbols = observed data; solid dots with connecting lines = predictions.

Table 1

Glossary of Parameters for the Discrete-Slots and Shared-Resources Models

Parameter symbol	Parameter description
Discrete-slots (DS) models	
Parameters common to all DS models	
m_i	memory-state probability at set-size i
v_C	mean correct drift rate on memory-based change accumulator
v_{NC}	mean correct drift rate on memory-based no-change accumulator
v_G	mean drift rate on guessing accumulators
$G_C(j)$	response-threshold parameter on guess-change accumulator at change-probability level j
$G_{NC}(j)$	response-threshold parameter on guess-no-change accumulator at change-probability level j
A_M	start-point variability on memory-based accumulators
A_G	start-point variability on guessing-based accumulators
σ_M	between-trial standard deviation of memory-based drift rate
σ_G	between-trial standard deviation of guessing-based drift rate
t_0	base time
DS1	
R_C	response-threshold parameter on memory-based change accumulator
R_{NC}	response-threshold parameter on memory-based no-change accumulator
DS2	
$R_C(i)$	response-threshold parameter on memory-based change accumulator at set-size level i
$R_{NC}(i)$	response-threshold parameter on memory-based no-change accumulator at set-size level i
DS3	
$ss_C(i) + cp_C(j)$	additive response threshold parameters on memory-based change accumulator at set-size level i and change-probability level j
$ss_{NC}(i) + cp_{NC}(j)$	additive response threshold parameters on memory-based no-change accumulator at set-size level i and change-probability level j
Continuous (C) models	
Parameters common to all C models	
$v_C(i)$	mean correct drift rate on change accumulator at set-size level i
$v_{NC}(i)$	mean correct drift rate on no-change accumulator at set-size level i
A	start-point variability parameter
σ	between-trial standard deviation of drift rate
t_0	base time
Model C1	
$R_C(j)$	response-threshold parameter on change accumulator at change-probability level j
$R_{NC}(j)$	response-threshold parameter on no-change accumulator at change-probability level j
Model C2	
$ss_C(i) + cp_C(j)$	additive response threshold parameters on change accumulator at set-size level i and change-probability level j
$ss_{NC}(i) + cp_{NC}(j)$	additive response threshold parameters on no-change accumulator at set-size level i and change-probability level j
Model C3	

Parameter symbol	Parameter description
$v_C'(i)$	mean error drift rate on change accumulator at set-size level i
$v_{NC}'(i)$	mean error drift rate on no-change accumulator at set-size level i

Table 2

BIC Fits of the Models to the Individual-Subject Data of Experiment 1

Subject	Model					
	DS1	DS2	DS3	C1	C2	C3
1	424	139	0	134	162	177
2	410	120	0	240	245	207
3	484	0	10	137	99	12
4	930	43	0	111	119	124
5	0	28	51	71	79	77
6	405	82	0	255	237	121
7	240	25	17	155	0	20
8	24	0	13	55	64	96

Note. For Subjects 1–8, the absolute BIC values of the best-fitting model were –5732, –4228, –4336, –3381, 1513, –7780, –1507, and –1319, respectively. Minimum BIC value for each subject is indicated in boldface. DS = discrete-slots; C = continuous; BIC = Bayesian information criterion.

Table 3

Best-Fitting Parameters From the Discrete-Slots Models for Subjects 3 and 1 of Experiment 1

Parameter	Subject, model	
	S3, DS2	S1, DS3
m_1, m_2, m_3	.960, .684, .431	.968, .860, .707
v_C, v_{NC}	1.306, 1.207	1.036, .828
v_G	.371	.509
$G_C(1), G_C(2), G_C(3)$	.142, .142, .096	.153, .127, .125
$G_{NC}(1), G_{NC}(2), G_{NC}(3)$	.083, .126, .150	.131, .155, .191
A_M, A_G	.263, .192	.062, .317
σ_M, σ_G	.222, [.100]	.225, [.100]
t_0	.100	.100
$R_C(1), R_C(2), R_C(3)$	.336, .456, .519	
$R_{NC}(1), R_{NC}(2), R_{NC}(3)$	.257, .319, .323	
$ss_C(1), ss_C(2), ss_C(3)$		[.000], .045, .065
$cp_C(1), cp_C(2), cp_C(3)$		.357, .322, .311
$ss_{NC}(1), ss_{NC}(2), ss_{NC}(3)$		[.000], .035, .051
$cp_{NC}(1), cp_{NC}(2), cp_{NC}(3)$		.220, .232, .267

Note. The R , ss , and cp parameters denote the extent to which the memory-based response thresholds are incremented beyond the value of the A_M start-point parameter. The G parameters denote the extent to which the guessing-based response thresholds are incremented beyond the value of the A_G start-point parameter. Parameters in brackets are held fixed at default values for scaling convenience. S = subject; DS = discrete-slots model; m = memory-state probability; v_C = mean correct drift rate on memory-based change accumulator; v_{NC} = mean correct drift rate on memory-based no-change accumulator; v_G = mean drift rate on guessing accumulators; G_C = response-threshold parameter on guess-change accumulator; G_{NC} = response-threshold parameter on guess-no-change accumulator; A_M = start-point variability on memory-based accumulators; A_G = start-point variability on guessing-based accumulators; σ_M = between-trial standard deviation of memory-based drift rate; σ_G = between-trial standard deviation of guessing-based drift rate; t_0 = base time; R_C = response-threshold parameter on memory-based change accumulator; R_{NC} = response-threshold parameter on memory-based no-change accumulator; ss_C = set-size threshold parameters on change accumulator; ss_{NC} = set-size threshold parameters on no-change accumulator; cp_C = change-probability threshold parameters on change accumulator; cp_{NC} = change-probability threshold parameters on no-change accumulator.

Table 4

p Values Associated With Kolmogorov-Smirnov Tests of Differences Between the RT Distributions at Set Sizes 5 and 8

Subject	Response type	Change-probability condition			
		.3	.5	.7	Pooled
False alarms (FAs) and misses					
1	FA	0.202	0.581	0.676	0.152
	Miss	0.531	0.630	0.602	0.573
2	FA	0.007	0.380	0.141	0.041
	Miss	0.032	0.854	0.368	0.161
3	FA	0.372	0.329	0.498	0.960
	Miss	0.570	0.248	0.328	0.855
4	FA	0.014	0.378	0.062	0.392
	Miss	0.218	0.055	0.389	0.146
5	FA	0.704	0.224	0.946	0.203
	Miss	0.336	0.975	0.292	0.137
6	FA	0.317	0.267	0.056	0.167
	Miss	0.452	0.402	0.094	0.993
7	FA	0.144	0.023	0.018	0.047
	Miss	0.160	0.622	0.136	0.031
8	FA	0.087	0.205	0.396	0.123
	Miss	0.815	0.856	0.370	0.720
Hits and correct rejections (CRs)					
1	Hit	0.058	0.002	0.001	0.000
	CR	0.042	0.001	0.219	0.000
2	Hit	0.007	0.031	0.001	0.000
	CR	0.028	0.033	0.463	0.029
3	Hit	0.009	0.000	0.000	0.000
	CR	0.629	0.380	0.005	0.130
4	Hit	0.001	0.000	0.000	0.000
	CR	0.006	0.000	0.000	0.000
5	Hit	0.037	0.731	0.051	0.004

Subject	Response type	Change-probability condition			
		.3	.5	.7	Pooled
6	CR	0.000	0.225	0.696	0.000
	Hit	0.031	0.048	0.012	0.000
	CR	0.000	0.000	0.019	0.000
7	Hit	0.000	0.001	0.000	0.000
	CR	0.000	0.006	0.000	0.000
	Hit	0.886	0.473	0.044	0.032
8	CR	0.064	0.034	0.955	0.021

Note. RT = response time.

Table 5

RGB-Values for the Stimuli Used in Experiment 2

Stimulus	Red	Green	Blue
black-1	0	0	0
black-2	55	55	55
white-1	243	243	243
white-2	200	200	200
red-1	254	0	0
red-2	170	0	0
blue-1	0	0	132
blue-2	68	110	255
green-1	0	255	0
green-2	0	193	0
yellow-1	250	255	67
yellow-2	255	211	0
orange-1	255	109	0
orange-2	255	151	0
cyan-1	0	255	255
cyan-2	0	190	232
purple-1	135	0	128
purple-2	107	0	54
dark-blue-green-1	0	89	100
dark-blue-green-2	66	134	133

Note. RGB = red– green– blue.

Table 6

BIC Fits of the Models to the Individual-Subject Data of Experiment 2

Subject	DS3	C2	H1	H2
1	26	107	16	0
2	27	65	16	0
3	19	0	11	1
4	31	0	41	27
5	20	297	12	0

Note. Minimum BIC value for each subject is indicated in boldface. For Subjects 1–5, the absolute BIC values of the best-fitting model were –342.5, –3,438.3, –1,971.6, 664.3, and 1,087.3, respectively. DS3 = Discrete-Slots 3; C2 = Continuous 2; H1 = Hybrid 1; H2 = Hybrid 2; BIC = Bayesian information criterion.