Published in final edited form as:

J Exp Psychol Learn Mem Cogn. 2016 May ; 42(5): 700–722. doi:10.1037/xlm0000197.

Reasoning and memory: People make varied use of the information available in working memory

Kyle O. Hardman and Nelson Cowan

University of Missouri

Abstract

Working memory (WM) is used for storing information in a highly-accessible state so that other mental processes, such as reasoning, can use that information. Some WM tasks require that participants not only store information, but also reason about that information in order to perform optimally on the task. In this study, we used visual WM tasks that had both storage and reasoning components in order to determine both how ideally people are able to reason about information in WM and if there is a relationship between information storage and reasoning. We developed novel psychological process models of the tasks that allowed us to estimate for each participant both how much information they had in WM and how efficiently they reasoned about that information. Our estimates of information use showed that participants are not all ideal information users or minimal information users, but rather that there are individual differences in the thoroughness of information use in our WM tasks. However, we found that our participants tended to be more ideal than minimal. One implication of this work is that in order to accurately estimate the amount of information in WM, it is important to also estimate how efficiently that information is used. This new analysis contributes to the theoretical premise that human rationality may be bounded by the complexity of task demands.

Introduction

Working memory (WM) researchers have long been interested in the amount of information people can actively store in mind in an accessible state (e.g. Cowan, 2001; Miller, 1956). The amount of information that can be held in WM has been correlated with important cognitive abilities (e.g. fluid intelligence; Colom, Abad, Quiroga, Shih, & Flores-Mendoza, 2008; Engle, Tuholski, Laughlin, & Conway, 1999). A less-studied topic is how well participants use the information they have in mind when performing WM tasks (Chen & Cowan, 2013). Some tasks require participants not only to maintain representations in WM, but also to reason about those memory contents. We will call this kind of reasoning about WM representations information use, and it can be added to other ongoing issues about the nature and use of working memory (e.g., Suchow, Fougner, Brady, & Alvarez, 2014).

The present research, which focuses on information use in WM tasks, can be viewed as a new approach to the issue of whether human thought processes are rational (as suggested,

for example, by Anderson, 1991; Danker, Gunn, & Anderson, 2008). Previous findings seem to suggest that humans are sometimes rational and sometimes not. An example is probability matching, which occurs in a two-choice guessing situation in which the payoff is the same for correctly choosing A or B on a trial, but A is more frequently correct. Although the rational decision would be to select A consistently, participants usually probability match, selecting A on a proportion of trials roughly commensurate with the proportion in which it occurs (e.g., Vulkan, 2000). Yet, Shanks, Tunney, and McCarthy (2002) found that with sufficient training, feedback, and incentives, participants moved toward the rational response strategy (in the example, always selecting Choice A).

Some researchers have demonstrated situations in which people do not respond rationally (Tversky & Kahneman, 1974, 1981). One reason for irrational behavior is that the ability to respond rationally is limited by the ability to calculate optimality. Thus, it has often been suggested that individuals respond according to a bounded rationality, in which the ability to process information is limited but the response is optimal given the proportion of the information that was processed (Gigerenzer & Goldstein, 1996; Kahneman, 2003). According to this line of argument, it should be revealing to examine reasoning within a WM task, in which memory capacity is loaded at the same time that a reasoning strategy might help performance. That is the case in the experiment providing the data base for the present theoretical investigation. We might anticipate that WM and reasoning depend on a common resource (Baddeley & Hitch, 1974; Kane et al., 2004; Süß, Oberauer, Wittman, Wilhelm, & Schulze, 2002) in which case rationality might be curtailed in some individuals more than in others and, within some individuals, on some trials but not others.

The thoroughness of information use is not directly an aspect of WM storage, but it does have a potential impact on estimates of the amount of information that can be stored. A psychological process model of a WM task might assume that participants in the task will use all of the information available to them as ideally as possible. However, if the participants are not completely ideal when responding, the model will underestimate the amount of information available to the participant. Due to this potential issue, it is important to try to estimate how ideally participants use the information available to them when performing the test phase of WM tasks. If we knew how well participants use information, we could improve the quality of estimates of the amount of information in WM. In turn, the separation of WM storage from WM information use would improve our understanding of the relation between WM tasks and other cognitive abilities.

We will begin by describing a prior study on which we based our method. We will discuss the analytical approach used by that study to examine information use in WM and the limitations of that approach. We will then describe the analytical and experimental methods we used and the results of our analyses.

Prior Research

This study is based on a recent study by Chen and Cowan (2013) that investigated information use in WM tasks. In that study, participants performed two WM tasks and the amount of information in WM was estimated for both tasks under differing assumptions of information use.

One of the tasks used by Chen and Cowan (2013) was a change-detection (CD) task, which is widely used to estimate the amount of information in WM (e.g. Cowan, 2001; Cowan, Blume, & Saults, 2013; Pashler, 1988; Rouder et al., 2008). A single trial of a standard visual CD task begins with the presentation of a sample array of visual objects which must be remembered for a short interval, after which the participant must determine if a presented memory probe contains some kind of change from the sample array. In Chen and Cowan's study, the stimuli were letters presented at different spatial locations, with no repetitions of letters within a given trial. In the CD task, the probe was a single letter presented at one of the locations. In Experiment 2 of their study, which is most comparable to our design, the only kind of change that was possible was that the letter might be presented at a location other than the one at which it was originally presented. The participant was to indicate whether the letter/location pairing that he or she was probed with had been presented in the sample array or if the probe was a recombination of a letter and a location that were not originally paired together.

The other task used by Chen and Cowan (2013) was what we will refer to as a feature-matching (FM) task. In a FM task, a participant is shown a number of stimuli which possess at least two features to remember. After a short delay, the participant is asked to match one feature from the sample array, the probe feature, to the other feature it was originally paired with. For the letter/location stimuli used by Chen and Cowan, the sample arrays used in the FM task were the same as in the CD task. There were two kinds of test in the FM task. In one kind of test, the participant was shown a single letter (the probe feature) and was asked to choose the location at which that letter had been presented from among several possible response options, each of which was a location that was used on that trial. In the other kind of test a single location was probed and the participant was asked to indicate which letter had been presented there originally. In both test situations, all of the possible responses were available: If the probe feature was a letter, any of the original locations could be selected and if the probe feature was a location, any of the letters could be selected. Because none of the letters or locations in the sample array could repeat, there was always exactly one correct response option.

Because CD tasks have been widely used and modeled, Chen and Cowan (2013) used estimates of the amount of information in WM that were derived from the CD task as benchmarks to which the FM task could be compared. Estimates of the amount of information in WM were derived from the FM task under two models which made different assumptions about information use. In both models for the FM task, it was assumed that if participants knew the binding between the probe feature and the other feature with which it had originally been presented, they would respond correctly. The models differed, however, in the assumption about the process that participants performed if they did not know the binding between the probe feature and the correct response option. The minimal responder model assumed that when participants did not know which response option was originally paired with the probe feature, they would pick a response at random. The ideal responder model instead made the assumption that when participants did not know which response option the probe feature went with originally, they would eliminate response options that they knew to be incorrect before guessing from the remaining response options. For example, assume a participant is presented with the letters A through D and remembers

where the letters A and B are, but does not remember C and D. If the participant is tested on C, which is not in mind, the participant does not know which response option is correct. However, because no features were allowed to repeat in the sample array, the participant is able to use their memory of the unprobed letters to eliminate the locations at which the letters A and B were presented. Once these 2 known-to-be-incorrect response options are eliminated, the participant guesses at random from the 2 remaining response options. The two models for the FM task result in different estimates for the amount of information in WM, because the more efficiently information is used, the less information is needed to perform the task at an equivalent level of accuracy.

Chen and Cowan (2013) estimated the amounts of information in WM in the two tasks under differing assumptions of information use, as described above. They found that the amount of information estimated to be held in WM in the benchmark CD task was quite similar to the amount of information that was estimated to be in WM in the minimal responder model for the FM task. Based on this, they concluded that information use in WM seems to be minimal. However, one important limitation of Chen and Cowan (2013) is that the benchmark model they used to estimate the amount of information in WM from the CD task was neither a minimal nor an ideal model, but an intermediate model that has not been validated with WM tasks in which bindings must be maintained. Thus, the benchmark model is not clearly the correct benchmark. There exist models which can be used as minimal or ideal responder models of a CD task in which bindings must be maintained. Two such models were described by Cowan et al. (2013) in Appendix A of that article. After we describe these minimal and ideal models, we will show how the model used by Chen and Cowan was an intermediate model.

One model, which we will refer to as the minimal model for the CD task, was called the “Reverse-Pashler” model by Cowan et al. (2013). Although those authors did not use this model for a case in which binding changes could occur, it is potentially applicable to that case. In the model, it is assumed that participants are able to recognize when the probe object is identical to one of the objects from the sample array as long as they have that object from the sample array in WM. If they fail to detect an object that they recognize, they will guess as to whether or not a recombination has occurred. To be explicit, participants cannot detect a recombination of features. We consider this to be a reasonable interpretation of what minimal information use might be: Although participants can successfully detect objects that match a WM representation, they do not or cannot perform the process required to detect a recombination.

The other model from Cowan et al. (2013), which we will call the ideal model for the CD task, assumes that participants not only detect recombinations, but do so in the most ideal way possible. To use the letter/location task for the example, when a recombination has occurred, the participant checks both ways in which he or she could detect a recombination. One way is that the participant remembers that at the tested location, the letter is not the same as was presented at that location originally. The other way is that the participant remembers that the letter that he or she is being tested on was presented at a different location originally. Thus, when there has been a recombination, a participant can detect it through knowledge of either of two bindings: The probed location and the letter that was

originally presented there or the probed letter and the location at which that letter was originally presented. In the ideal CD model, as in the minimal model, participants are assumed always to detect a probe that is identical to an object from the sample array if that object is in WM.

In the intermediate CD model used by Chen and Cowan (2013), participants are assumed to be able to detect a recombination through one, but not both, of the features of the probe object. In an attempt to detect a recombination, a participant might check to see if he or she remembers what letter was originally presented at the probed location. If this information is lacking, the participant does not also check to see if he or she remembers the location at which the probed letter was originally presented. Thus, the model used by Chen and Cowan can be roughly understood as being about halfway in between the minimal and ideal models – an intermediate model. Again, just as in the other CD models described above, participants were assumed to detect probes that were identical to a sampled object. Note that this intermediate model for the CD task is the widely used model suggested by Cowan (2001), but that the model was not originally suggested for use with bindings. Before a mathematical model is used for a new type of task or for substantially different stimuli, the specification of the model should be carefully considered in order to determine what exactly the model says about the process it assumes that participants will use in the new situation (Rouder, Morey, Morey, & Cowan, 2011).

Research Strategy

Experimental Method Overview—To address the question of how efficiently information is used in WM tasks, we will follow the basic approach of Chen and Cowan (2013), but with some modifications. We will use the same two types of task that were used by Chen and Cowan, but with different stimuli. Our stimuli are visual objects that possess a color and an orientation as the two features, as can be seen in Figure 1. The reasons for changing the stimuli from what were used by Chen and Cowan were (1) to eliminate automatically verbalized stimuli (letters), and (2) to eliminate location as a feature, given that location has been shown to be special compared to other features (Treisman & Zhang, 2006). The result of these changes are that the two features are more similar to one another, as they are both non-location visual features.

For both the CD and FM tasks, on each trial a sample array of 2 to 5 objects is presented. For the colored and oriented objects in each sample array, colors and orientations were not allowed to repeat. That characteristic of the experiment made it possible for participants to use deductive reasoning in the test situations, as will be described in our modeling approach. In the CD task test situation, participants had to judge whether the color and orientation of a single probe object were both drawn from a single object from the sample array (an unrecombined probe) or were drawn from two different objects (a recombined probe). In the FM task test situation, participants were shown a single probe feature, which was either a color or an orientation, and were tasked with selecting the feature of the other type that was originally presented with that probe feature. For example, following a 4-item array, if the probe was one of the colors then the participant would have to select the appropriate orientation that had been combined with that color, from the 4 orientations that had been

presented in the sample array. The data from these tasks will serve as the basis for a mathematical modeling analysis.

Modeling Approach—The purpose of our modeling approach is to estimate the values of a parameters that indicate how much of the information in WM is actually used, on a continuum from minimal to ideal. The minimal, intermediate, and ideal models we have described all make categorical assumptions about information use (e.g. participants are fully minimal or fully ideal). However, it is not necessary to make categorical assumptions about information use in order to use mathematical modeling. Instead, our primary analysis will use mixture models of both the CD and FM tasks. In the mixture models, a freely-estimated parameter controls the probability that the participant will enter a minimal responder state or an ideal responder state on each trial. In this way, the mixture models allow each participant to be completely ideal, completely minimal, or any combination of the two states. This approach is appealing due to the fact that it avoids making the assumption that information use is all-or-none. However, it is relatively complex to implement and it is not obvious that other, simpler approaches are insufficient to study information use in WM. For this reason, we will argue against two other approaches that – while much simpler to implement – we believe are insufficient for studying information use in WM.

The first approach that we find insufficient is the use of benchmark models for the given problem. With the two new ideal and minimal models for the CD task available to us, we could repeat the analysis of Chen and Cowan (2013) with the two new models as the benchmark models. However, the existence of more than one benchmark model immediately makes the use of benchmark models dubious because it isn't possible for two different models which make diametrically opposed assumptions about psychological processes to both be correct. The use of a benchmark as a measure of the correctness of a different model necessitates that the benchmark actually be known to be correct, not just one of many possible options. With the existence of multiple possible benchmark models, none of which is known to be correct, we must find an approach that does not use benchmarks in order to learn about information use in WM.

A different approach we could use would be to use the agreement of ideal or minimal models from the two tasks to determine which models were correct. To do this, we could perform an analysis like that of Chen and Cowan (2013) with the addition of the new ideal and minimal CD models. By comparing the WM capacity estimates of the minimal and ideal models for the CD task with the WM capacity estimates of the same models for the FM task we could look for agreements in WM capacity estimates under the same assumptions of information use. Let us assume that we did this analysis and that we found that the ideal model from the CD task agreed with the ideal model from the FM task, but the minimal models for the two tasks did not agree. Does this tell us that the ideal models are both correct? Not really, because there could be an intermediate point somewhere between completely ideal information use and completely minimal information use at which the models could also agree. For example, participants could be 70% of the way from minimal to ideal and the two models might still agree even though both assumed completely ideal information use. One could even hypothesize that participants are more ideal in one task than the other, creating even more ways in which the models could agree in terms of WM

capacity estimates for the same set of data. In order to use simultaneous validation of models, both models must only be able to agree at one point for a given set of data. The evidence provided by a simultaneous validation strategy is weak in this case because there are many ways in which the models might agree, not just one.

These simpler approaches are not able to give us unambiguous answers about information use in WM. For this reason, we will use a modeling approach in which we allow the proportion of trials on which information was used ideally, as opposed to minimally, to be freely estimated for each participant. We will estimate how ideally information is used by fitting the model parameters to the data, estimating both how much information participants have in WM and how well they use that information. This approach avoids a need for benchmarks because information use for each task will be estimated independently of the other task. Also, because each task is treated independently, there is no need for us to try to simultaneously validate one task's model based on the results of the other task. Additionally, because we will estimate how ideally each participant uses information in WM, we avoid problems that may arise from aggregation of data (c.f. Brown & Heathcote, 2003; Estes, 1956; Pratte, Rouder, & Morey, 2010). For example, if half of the participants were minimal and the other half were ideal, we might mistakenly conclude that participants are moderate information users on the basis of aggregate data when in fact all participants are extreme, just at different extremes. We will now turn to a full description of the mathematical models we used.

Model Descriptions

For our thorough description of the models we used, we will begin with the model for the FM task. For reference, see Table 1 for brief definitions of the parameters used in the models, along with definitions of various terms related to the tasks. A multinomial process tree for the FM model is shown in Figure 2, providing a graphical representation of the process of memory and decision making that we assume is used by participants in the task. The FM model accounts for the possibility that participants' attention may lapse on some trials with an attention parameter, A . The use of the attention parameter is warranted to explain why participants occasionally make errors even at array sizes that are less than K (Rouder et al., 2011). With probability $1 - A$, the participant is not paying attention and the model assumes that he or she has no information in WM and must choose a response option at random. There are N objects in the sample array and N corresponding response options, exactly one of which is the correct response, so the participant guesses correctly with probability $1/N$. To summarize the model in the case of a participant who is paying attention, the participant begins by attempting to detect the correct response and, if that fails due to a lack of the necessary information in WM, enters a guessing state in which a mixture parameter controls the probability of ideal or minimal information use.

In more detail, the FM model assumes that a participant has K objects in WM for which the color/orientation binding is known which were selected from the N total objects in the sample array. If the participant has the object that contained the probe feature in WM, which happens with probability $\min(K/N, 1)$, he or she is able to answer correctly because the response option that went with the probe feature in the sample array is known by the

participant. Note that we will use multiple array sizes and that a single K value will be estimated based on all array sizes. As such, it is possible for K to be greater than N at smaller array sizes. If K is greater than N , then the number of objects in WM will be set equal to N for that array size. This makes sense under the assumption that if a participant has a WM capacity of K items but is presented with fewer than K items in the sample array, they will encode all N of those items, not completely filling their WM capacity. By taking the $\min$ of K/N and 1, we are stating that, even if $K > N$, the participant does not know the correct answer with probability greater than 1 because only N items were actually encoded.

If the participant does not know the correct answer, which happens with probability $1 - \min(K/N, 1)$, he or she enters a guessing state. In the guessing state, with probability β , the participant enters an ideal responding state in which he or she eliminates $\min(K, N - 1)$ incorrect response options before guessing from the remaining options. The minimum of K and $N - 1$ is used because there are only $N - 1$ incorrect response options that can be eliminated and because K is the asymptotic WM capacity, it can be greater than $N - 1$ when N is small. With probability $1 - \beta$, the participant enters a minimal responding state in which he or she picks a response option at random, guessing correctly with probability $1/N$. With $A = 1$ and β set to either 0 or 1, our model for the FM task is the same as the minimal-responder ($\beta = 0$) or ideal-responder model ($\beta = 1$) used by Chen and Cowan (2013).

The model for the CD task follows the same principles as the model for the FM task. The CD model is shown in two parts in Figures 3 and 4, which correspond to trials on which there was a recombination of features and trials on which there was no recombination of features, respectively. In the CD task, a response was recorded as a hit if the participant correctly responded that the probe object was a recombination and coded as a correct rejection if the participant correctly responded that the probe object was an unrecombined object, which follows the typical coding used in change-detection tasks since Pashler (1988). Like the FM model, the CD model has an attention parameter A . Regardless of whether or not there has been a recombination, with probability $1 - A$ the participant is not paying attention, has no information in WM, and guesses that there has been a recombination with probability u , where u is the uninformed guessing rate (because the participant was not paying attention, he or she has no information to use to help them guess, hence “uninformed” guessing). If the participant is paying attention, the model differs depending on whether there actually has or has not been a recombination.

If there actually has been a recombination of features in the CD task, the model follows Figure 3. With probability $1 - \alpha$, the participant is using information minimally and is incapable of detecting recombinations. In this case, the participant guesses that there has been a recombination with probability g_M . However, with probability α the participant is using information ideally and can detect recombinations. In this case, the participant has two chances to know the object from which the probe object’s features were drawn. As an example, assume that the probe object is orange and has an orientation of 45° and that the participant performing the task knows that each feature can only repeat once in each sample array, as they should. In that case, if the participant remembers that the color orange went with an orientation different from 45° originally, he or she knows that there has been a recombination. Failing that, if the participant knows that the orientation 45° originally went

with a color other than orange, he or she knows that there has been a recombination. Without knowledge of the original binding of either feature of the probe object, the participant must guess about whether or not there has been a recombination. More formally, with probability $\min(K/N, 1)$, the participant knows the object from which the first feature of the object (color in the above example) was drawn. Because the participant knows the object from the sample array that possessed this first feature (color) possessed a different second feature (orientation) than the probe object, he or she is able to correctly respond that there has been a recombination. If the participant does not detect that the first feature was originally paired with a different second feature, he or she can still detect that the second feature of the probe object was originally paired with a different first feature if the participant remembers that the object that possessed the second feature of the probe object. This second detection happens with probability $\min(K/(N-1), 1)$ ¹. To finish with the branches of the model, if the participant did not detect that a recombination occurred on the basis of the two features of the object, he or she would guess that a recombination occurred with probability g_M .

In the CD model, three different guessing rates are used, depending on whether the participant is guessing due to inattention (u), due to responding minimally (g_M), or due to responding ideally but failing to detect a recombination (g_I). As discussed further in Appendix A, g_M and g_I reflect different kinds of informed guessing, which, in contrast to uninformed guessing (u), depend on what information a participant has in WM and how ideally he or she has processed that information. Thus, g_I is the probability that a participant guesses that there has been a recombination given that he or she has $\min(K, N)$ objects in WM and that he or she used that information ideally when attempting to detect recombinations. g_M is defined analogously for the case in which information was used minimally to detect recombinations. Although there are three guessing rates, there are only two free parameters in the model that are related to guessing: u , the uninformed guessing rate, and P_{IG} , which can be read as “Probability of using Informed Guessing”. u is what the participant believes is the probability that a probe object will be recombined in the absence of any other information. If a participant has no information about the sample array (i.e. the participant was not paying attention), he or she will guess that there has been a recombination with probability u . Turning to the other parameter related to guessing, P_{IG} , it is reasonable to think that using informed guessing would demand some cognitive resources and that participants might not always use informed guessing perfectly, instead falling back on the less-demanding uninformed guessing. To accomplish this, P_{IG} mixes between perfect informed guessing behavior and uninformed guessing behavior. Thus, g_M and g_I are defined as follows

¹We make the assumption that K objects are always encoded on each trial, assuming that $K < N$. The denominator of $N-1$ is used because in order for a participant to check if he or she knows the object from which the second feature of the probe object was drawn, the participant must not have known the object from which the first feature was drawn, otherwise he or she would have detected the recombination based on that knowledge of the first feature and would not be performing the check of the second feature. Additionally, because the probe object is recombined, its features have been drawn from two different sample objects. Because the participant has K objects in mind but does not know the object from which the first probe feature was drawn, he or she must not know one of the sample objects. This sample object that is not known cannot be the one from which the second feature was drawn, so their probability of knowing the object from which the second feature was drawn increases because the K objects known to the participant know are drawn from $N-1$ objects that include the object with the second feature of the probe object.

$$g_M = P_{IG} * \gamma_M + (1 - P_{IG})u \quad (1)$$

$$g_I = P_{IG} * \gamma_I + (1 - P_{IG})u \quad (2)$$

where γ_M and γ_I are the informed guessing rates that a participant would use if he or she were using informed guessing perfectly, the equations for which are given in Appendix A. Thus, P_{IG} reflects yet another kind of information use that is distinct from α . α is related to detecting recombinations and P_{IG} is related to the use of informed guessing.

The other type of test in the CD task is one in which there has not been a recombination of features for the probe object (i.e. the probe object is a whole object from the sample array). In this case, the CD model follows Figure 4. With probability $\min(K/N, 1)$, the participant has the probe object in WM and correctly responds that he or she has seen it before. However, if the participant does not have the probe object in WM, he or she must guess. The participant's guessing rate depends on whether or not he or she used information ideally on that trial (i.e. it depends on α). If the participant used information ideally, he or she guesses that a recombination occurred with probability g_I . If the participant did not attempt to detect recombinations, he or she guesses that a recombination occurred with probability g_M . If the probe object is not made from a recombination of features, the ability of a participant to detect recombinations clearly does not depend on whether or not that participant attempted to detect recombinations, because it is impossible to detect a recombination if it has not actually occurred. The guessing rates for the CD model are the same regardless of whether there actually was or was not a recombination, because guessing by definition only occurs if the true state of the probe object is unknown and so cannot depend on the true state.

It should be noted that information use, as reflected by α and β , is defined in analogous ways for the models of both the CD and FM tasks. Both models define ideal information use as using deductive reasoning at test about the information in WM. For a recombined probe in the CD model, the participant knows that the object is not one of the whole objects that he or she has in WM. The participant can detect whether or not the probe is a recombination of features by determining for each feature if that feature went with a different feature in the sample array. If so, the participant knows that the probe is a recombination. The rate at which the participant does this checking is reflected by α . In the FM model, if the participant does not know the object that the probe feature came from, he or she cannot know the correct response option. However, the participant can check for each response option if it went with a feature other than the probe feature in the sample array, and, if so, eliminate it from consideration. The rate at which a participant does this checking is reflected by β . Thus, information use is defined very similarly for α (in the CD task) and β (in the FM task). Note that both α and β are unrelated to recognizing whole objects: In the CD task, α is not involved in detecting unrecombined probes and in the FM task, β is not related to selecting the correct response option when the probe feature comes from an object in WM. Our definitions of information use require that participants do something with the information beyond basic recognition processes. Although informed guessing, controlled by P_{IG} , in the CD task is a type of information use, it is distinct from the type of information

use reflected by α and β . While α and β are related to deductive reasoning about the concrete features of objects, P_{IG} is related to abstract, probabilistic reasoning about the amount of information in WM, how ideally that information was used to detect recombinations, and the probability that different types of probe objects would have gone undetected.

Our models are related to the class of sophisticated guessing models (e.g. Broadbent, 1967; Johnson, 1978; Prizmetal & Lyon, 1996), in which participants use partial information to narrow the range of possible responses. Sophisticated guessing models have typically been used to explain behavior in tasks that require participants to identify the letters in briefly presented words. For example, in a task in which participants are shown English words and must identify the letters in those words, a participant sees the letters “PAR_” and must fill in the last letter to complete the word. The participant could guess at random, selecting from the 26 letters in the alphabet, or the participant could use their knowledge of English words to limit the possible responses to the letters E, K, or T. In our tasks, sophisticated guessing is most applicable to guessing in the FM task, in which case participants could eliminate incorrect responses before guessing from the remaining options. In a general way, sophisticated guessing is also a good description of informed guessing in the CD task, where participants use knowledge about what they would have detected, given what they know, to guess more effectively.

Method

Participants

For this experiment, the data from 52 participants (32 female) with mean age 24.8 (standard deviation 10.7) years was included in the final analysis. Of the participants, 38 were recruited from an introductory psychology class at the University of Missouri and given partial course credit for their participation and 14 were recruited from the general population in Columbia, MO and were paid \$15 for their participation. Two additional participants took part in the experiment but were excluded from the final analysis. One participant was excluded for performing at chance on both tasks and one participant withdrew from the experiment before completing both tasks. All 52 included participants performed both the FM and CD tasks in their experimental session. This study was approved by the Campus Institutional Review Board of the University of Missouri –Columbia.

Materials

The stimuli used in both tasks were triangles that each had a color and an orientation. Examples of the stimuli can be seen in Figure 1. Eight highly-discriminable colors and eight orientations beginning at 0° and separated by 45° were used. The approximate color names and exact red, green, and blue (RGB) coordinates for the object colors were red (229, 0, 0), orange (229, 140, 0), yellow (255, 229, 0), green (0, 229, 57), cyan (0, 255, 255), blue (25, 82, 255), dark violet (130, 26, 242), and pink (229, 0, 172). The background color was a dark grey (40, 40, 40) and text and object outlines were white (255, 255, 255). The triangles had two internal angles of 70° and one internal angle of 40°, creating two dull points and one sharp point. The orientation the triangle pointed to was based on the direction pointed to by the sharp point. The triangles were 1° of visual angle long from the end of the sharp point

to the middle of the opposite side (between the dull points). In each sample array, between 2 and 5 objects were presented in an invisible circle with a radius of 6.4° of visual angle around the center of the screen. The circle on which the objects were placed passed through the center of an imaginary, smaller circle surrounding each triangle and touching each of its three points. In the sample arrays, colors and orientations were prevented from repeating so that participants would be able to reason deductively about unknown objects given knowledge about other objects. The participants were not explicitly informed that this was the case; we assumed that participants would recognize this fact after a brief exposure to the stimuli as it was very obvious.

Procedure

For each of the CD and FM tasks, participants completed a trial block of 240 trials with a break in the middle of the trial block (there were no practice trials for either task). Within each task's trial block, the array sizes (2 to 5 objects) were randomly intermixed. At the beginning of the trial block for each task, participants were given instructions for the coming task with images and text describing the task. The order in which participants received the tasks was counterbalanced across participants. No effects of counterbalancing order were observed. For both tasks, each trial began with a fixation cross for 1000 ms, followed by a 500 ms blank. Then the sample array was presented for 500 ms followed by a 1000 ms maintenance period. Following maintenance, participants were tested in a task-dependent way.

In the test phase of the CD task, a single probe object containing both a color and an orientation was presented in the center of the screen. The probe object was either identical to an object that had been presented in the sample array (an unrecombined probe) or had a color drawn from one object in the sample array and an orientation drawn from another object (a recombined probe). The participant responded by pressing a key on a standard American keyboard. They pressed 'S' to indicate that the probe object was the same as an object from the sample array or 'D' to indicate that the probe object was different from any object from the sample array (i.e. that it was a recombination probe). Participants were explicitly instructed that the only possible change would be the recombination of features from different objects. In the CD task, due in part to a programming error, there were 64 trials at array sizes 2 and 3 and 56 trials at array sizes 4 and 5. Half of the trials had a recombined probe and the other half had an unrecombined probe.

In the test phase of the FM task, the probe feature, which was a feature from one of the objects from the sample array, was presented in the middle of the screen. The response options were presented in a horizontal line 6.4° below the probe feature. The response options were all of the features from the sample array that could possibly be paired with the probe feature. For example, if the probe feature was a color, the response options would be the orientations of each of the objects from the sample array. In order to remove color information from the orientations at test, orientations were indicated by triangles outlined in white and filled with the background color (white was not one of the colors used for the stimuli). Similarly, colors were stripped of orientation information by presenting the color as a filled circle of that color. The response options were numbered from left to right with the

numbers increasing from 1 presented in an 18 point font. The participant was to select the correct response option from the list by pressing the keyboard key of the same number from the number row of a standard US keyboard (number pad responses were disallowed). In the FM task, there were 60 trials at each array size, half of which had color as the probe feature and half of which had orientation as the probe feature.

In both tasks, the location of probes did not correspond to locations of the stimuli in the sample array. This is important because if location could be used to make responses, participants would be able to use bindings that included location information, while we wanted participants to only be able to use color/orientation bindings. To remove location information, we substantially varied the way in which stimuli were presented between sample and test, as can be seen in Figure 1. In both tasks the sample stimuli were presented in a ring. In the CD task, the probe object was in the center of the ring defined by stimuli in the sample array, so there was no information that could link the probe object to any of the sample locations. In the FM task, the probe feature was also presented in the center of former ring. The response options were presented in such a way that it was possible for locations of at least some of the response options to overlap locations at which sample stimuli had been presented. However, given the random shuffling of the response options it would be impossible for participants to use location cues to help make their response. In addition, due to the fact that the sample stimuli were presented in a ring and the FM response options were presented in a horizontal line, it seems unlikely that participants would try to use location cues to make their response.

Results

Before the data were subjected to analysis, trials with response times longer than 20 s were removed, which resulted in a loss of only 0.16% of the trials. Summary statistics from the cleaned data are presented in Table 2. These data are also plotted in Figure 5, but model predictions are also plotted in that figure, making the raw data harder to see. The data show the expected pattern, in that performance declined as array size increased. In addition, the FM appeared to be more difficult than the CD task, with both lower accuracy and higher response times, as can be seen in Table 2. Lower accuracy is to be expected as chance performance on the FM task is $1/N$ whereas chance performance on the CD task is always $1/2$. However, the higher response times suggest that the FM task is more difficult than the CD task, which is discussed further.

In the FM task, because there are two different features that are used as the probe feature, the FM task could be conceptualized as two different tasks: The color probe FM task and the orientation probe FM task. In order to determine whether to model the two probe feature types in the FM task as separate tasks, we analyzed proportion correct in the FM task as a function of array size by probe feature using a Bayesian ANOVA from the BayesFactor package (R. D. Morey & Rouder, 2015) for R (R Core Team, 2015). We found evidence that there was no main effect of probe feature, $BF_{01} = 4.0$, and no interaction between probe feature and array size, $BF_{01} = 5.6$. As we used Bayesian analytical techniques in this study, we will report the result of statistical tests with Bayes factors, using *BF* as a shorthand for Bayes factor. When given, the subscripted numbers following *BF* indicate the numerator and

denominator models for the Bayes factor, with 0 representing the null model and 1 representing the alternative model. Because we did not find any effects of probe feature in the FM task, we combined data from the color and orientation probe trials in all further analyses.

Closed-form estimation of K

In order to allow easier cross-study comparison, we initially analyzed our data in the same way as Chen and Cowan (2013), with the addition of the minimal and ideal CD models. This involved using closed-form solutions for K , the number of items in WM, at each array size under different assumptions of information use. The closed-form solutions for both the minimal and ideal FM models, and the intermediate CD model, are provided by Chen and Cowan. The minimal CD model solution is provided in Appendix A of Cowan et al. (2013) under the heading “Model 2”. The ideal CD model solution is provided in Appendix B of Zhang, Xuan, Fu, and Pylyshyn (2010). As discussed in the introduction, none of these models can be considered to be a benchmark model. Agreement between two models should not be interpreted as indicating that those models are both correct. The result of this analysis is presented in Figure 6. The first thing to notice is that our results are very different from the results of Chen and Cowan, primarily in that neither FM model agrees with the intermediate CD model. The main point to take away from this figure is that it tells us very little about information use in WM. There is no discernible pattern in the K estimates that would allow us to make any conclusions about how ideally information is used in WM, even if we could assume that one of the models is a benchmark, which we cannot. Thus, we will proceed to our primary modeling approach with freely-estimated continuous information use parameters.

Parameter Estimation

All of the parameters in the CD and FM models cannot be estimated at each array size because there would be more free parameters than data points at each array size. To simplify the model in a principled manner, each participant’s parameters will be fixed across all array sizes. The K parameter, which represents WM capacity, is theoretically a stable characteristic of each participant that should not change depending on array size. However, each participant was given a different K parameter for each task to account for potential differences in task difficulty suggested by the raw data, which is discussed further below. The attention parameter, A , should not vary depending on array size because inattention should occur equally often at each array size. We do not know whether α , β , and P_{IG} are the same at all array sizes. It is possible that participants change their information use strategies as the memory load increases at higher array sizes. We chose to treat the information use parameters as invariant with respect to array size to control model flexibility. If information use varies with array size, even with information use treated as invariant we should still be able to estimate the mean level of information use.

The parameters of the model were estimated using a Bayesian approach, the details of which are discussed in Appendix B. In brief, hierarchical priors were placed on both K parameters, A , and u in order to provide constraint on the estimation of those parameters. Uniform priors on the interval $[0, 1]$ were placed on α , β , and P_{IG} . The use of uniform priors on the

information use parameters is related the method we used to categorize participants as minimal, moderate, or ideal information users, which is discussed in a later section. The detailed specifications of the priors we use on the parameters are given in Appendix B.

Model Fit

We can examine the fit of the models by considering the posterior predictive distributions from the model. A posterior predictive distribution gives the density for a new, as yet unobserved observation given the posterior distributions of the model parameters. There is a posterior predictive distribution for each data point, where a data point is defined as, for example, the hit rate at array size 3 for participant 12. The average prediction of the model for a single data point can be obtained by taking the mean of the posterior predictive distribution for that data point. We plotted these posterior predictive means along with the actual data in Figure 5. As can be seen, there are some group-level biases in the models, such as the way in which the FM model slightly overpredicts proportion of correct responses at array size 5 in that task. Note that for the CD task the hits and false alarms must be considered together when considering bias. However, in total, the group-level model predictions track the actual data fairly well at the group level, suggesting that the models are plausible.

The plots in Figure 5 are best suited to showing group-level bias in the models and they do not do as good of a job at showing variability at the level of individual data points. As suggested by Gelfand (1996), we can check that an appropriate proportion of the data falls into a central interval of the posterior predictive distributions. If too low a proportion of the data falls into the central interval and instead falls into the tails, then the model is not fitting the data well; if too great a proportion of the data falls into the central interval, the model is overfitting the data. We used both the central 80% and central 95% of the posterior predictive distributions and calculated the proportion of the data points that fell within these intervals. We found that 92.3% of the data fell into the central 80% of the posterior predictive distribution and that 97.7% of the data fell into the central 95%. Thus, there is some evidence that the model is overfitting the data, but not enough to be problematic. Based on our model fit diagnostics, there is no clear reason to doubt the validity of our models, so we will proceed to interpreting the parameter estimates.

Parameter Estimation Results

We start with a basic summary of the results for each parameter. Rather than working with the complete posterior distribution for each parameter for each participant, we will use the mean of the posterior distribution for each parameter as a point estimate of the value of that parameter for that participant. As a result, each participant has a posterior mean for each parameter and the summary results will be based on these posterior means. Summary statistics for the parameters are given in Table 3 and are also plotted in Figure 7. Parameter correlations and associated Bayes factors were estimated using the BayesMed package (Nuijten, Wetzels, Matzke, Dolan, & Wagenmakers, 2015) for R and are given in Table 4. The primary results of interest are discussed in the following sections.

WM Capacity Estimates—We observed mean WM capacity estimates of 1.54 in the FM task and 2.09 in the CD task. These values might at first seem to be too low given the commonly-cited WM capacity of 3 or 4 objects (e.g. Cowan, 2001), but it should be remembered that a capacity of 3 or 4 objects is typically found for simple stimuli that have only a single relevant feature. For our stimuli, the binding between two non-location features must be remembered and our K estimates are similar to those found by other studies which use similar stimuli (e.g. Cowan et al., 2013). Additionally, the K estimates from the two tasks were positively correlated, $r = .44$, $BF_{10} = 28.5$, which helps support the idea that the K estimates we obtain are reasonable. Another possible concern is that our K estimates were not the same for both tasks, but we do not think that this is a problem for the validity of our models, as we will argue in the Discussion.

Correlations between the information use parameters—As we have described, α and β , while from different tasks, are logically related to similar kinds of information use involving deductive reasoning about the features of objects. P_{IG} , on the other hand, is somewhat distinct from α and β in that it involves more abstract reasoning about the amount of information in WM. As a result, our primary prediction is that α and β would be positively correlated. In addition, there may be other positive correlations that are due to the fact that all three information use parameters could be thought to be related to a general reasoning ability.

In contrast to our predictions, we found that α and β were not correlated, $r = 0.10$, $BF_{10} = 0.14$. Although on the basis of the logic of our models, α and β represent similar abilities, it could be that the CD and FM tasks have test situations that require substantially different abilities and that the logical relationship between α and β does not hold up in reality. One important difference between information use in the CD and FM tasks is the coherence of the objects in the test display. In the CD task, a single object possessing both a color and an orientation must be compared to the objects in WM. In the FM task, multiple objects each possessing only a single feature must be processed. In addition, much more information is available at test in the FM task, which could change participants' strategies. These differences between the tasks could possibly explain the lack of a correlation between α and β . However, it still might be expected that there would be a positive correlation between α and β that is due to individual differences in, e.g., motivation.

Additionally, we found that α and P_{IG} from the CD task were correlated, $r = 0.45$, $BF_{10} = 35.9$. This correlation might result from participants tending to use information in a similarly ideal way for both of the kinds of information use reflected by α and P_{IG} in the CD task. Finally, we found that β and P_{IG} were not correlated, $r = 0.08$, $BF_{10} = 0.13$, as we might expect given that β and P_{IG} represent logically distinct kinds of information use and are from different tasks.

Correlations between WM capacity and information use—We are also interested in correlations between WM capacity and information use. We would hypothesize that participants with higher WM capacities are able to use information more efficiently. This hypothesis is based on research that suggests that both WM capacity and reasoning ability are related to a common ability, such as attention (e.g. Kane et al., 2004; Süß et al., 2002).

We found a strong positive correlation between the CD K and α , $r = 0.75$, $BF_{10} = 1.74 * 10^8$. The correlation between the CD K and P_{IG} was 0.29, but the evidence was inconclusive about whether the correlation was nonzero, $BF_{10} = 1.14$. Finally, the correlation between the FM K and β was 0.33, with some small evidence that the correlation was nonzero, $BF_{10} = 2.54$. Thus, there is some evidence that for the types of information use represented by α and β , that information use positively correlates with WM capacity.

Categorizing Information Use

Now that we have examined relationships between the main parameters of interest, we want to categorize the information use of participants as minimal, moderate, or ideal. The approach we used takes advantage of the fact that we have a posterior distribution for each parameter in our model. As an exemplar, we will focus on the posterior distribution of the α parameter of one of our participants, shown in Panel A of Figure 8. As can be seen, the mode of the posterior distribution is near 0.8 and most of the density is above 0.5. This would suggest that this participant is reasonably ideal. As discussed in the Parameter Estimation section, we used a uniform prior on α and naturally treat that prior as a baseline of belief about α . In Panel A of Figure 8, the density of the prior distribution is shown with the horizontal line. In standard Bayesian fashion, we will use the differences between the prior and posterior distributions to update our beliefs about how ideal this participant is for the kind of information use represented by α .

We will define intervals of the prior and posterior distributions of α that we will define as corresponding to minimal, moderate, and ideal information use. The minimal interval was from 0 to 0.2, the moderate interval was from 0.2 to 0.8, and the ideal interval was from 0.8 to 1. The proportion of the prior and posterior distributions that are in these intervals can be used to categorize each participant for each parameter. This is done by using the fact that a Bayes factor can be expressed as the ratio of prior and posterior odds ratios, as in

$$\frac{Post(C)}{Post(\bar{C})} = BF_C \frac{Prior(C)}{Prior(\bar{C})} \quad (3)$$

where BF_C is the Bayes factor for the hypothesis that the participant is in category C versus the hypothesis that the participant is not in category C , $Post(C)$ and $Post(\bar{C})$ are the posterior probabilities for and against the hypothesis (the bar over a letter indicates logical negation), and $Prior(C)$ and $Prior(\bar{C})$ are the prior probabilities for and against the hypothesis. $Post(C)$ is simply the proportion of the posterior distribution that is within the interval of category C and $Post(\bar{C})$ is $1 - Post(C)$. Similarly, $Prior(C)$ is the proportion of the prior distribution within the interval of category C , which is just the width of the category because the prior distribution has a height of 1 everywhere, and $Prior(\bar{C})$ is $1 - Prior(C)$. Thus, BF_C can be calculated with

$$BF_C = \frac{Post(C)}{1 - Post(C)} * \frac{1 - Prior(C)}{Prior(C)}. \quad (4)$$

The category for each participant by parameter combination was chosen on the basis of which of category had the greatest BF_C . We performed this categorization separately for α , β , and P_{IG} , making it possible that a participant could be categorized as being, e.g., minimal for α but ideal for β .

As an example of the categorization procedure, we will continue with the α parameter shown in Panel A of Figure 8. Suppose that we want to test the hypothesis that this α reflects ideal information use. The posterior mass within the ideal interval (0.8 to 1) is 0.39 and the prior mass within the same interval is 0.2. Thus, BF_C for the ideal category is

$$BF_C = \frac{0.39}{1-0.39} * \frac{1-0.2}{0.2} = 0.64 * 4 = 2.56.$$

Note that, because this is just an example, all values were rounded to 2 decimal places before each calculation. The BF_C for the minimal and moderate categories were 0.08 and 0.95, respectively. As such, the participant was categorized as ideal for α by this procedure. The winning Bayes factor was not strong, being less than 3, but this is not surprising because it is based on only a single participant's data. Across all participants, the mean (median) winning Bayes factors were 3.62 (2.60) for α , 2.58 (1.67) for β , and 2.41 (1.79) for P_{IG} . This suggests that the winning categorizations, while not typically definitive, were strong enough to consider useful given that the categorization takes place at the individual participant level.

Once we had categorized our participants, we also wanted to test whether the distribution of participants into categories was uniform. In the categorization procedure, participants have an equal chance of being categorized into each of the categories, because the evidence can equally support any categorization. Thus, in order to make the claim that our participants were categorized as, e.g., ideal more or less often than would be expected if the distribution of categories of information use were uniform, the proportion of participants categorized as ideal must be different from 1/3. For each of the categorizations for each parameter, we tested the hypothesis that the proportion of participants given that categorization was 1/3 using the “proportionBF” function of the BayesFactor package.

The results of the categorization procedure and the tests of proportions are plotted in Panels B and C of Figure 8. We found that participants were categorized as ideal more often than one third of the time for β , $BF_{10} = 2.0 * 10^3$, and P_{IG} , $BF_{10} = 9.1$. In addition, we found that participants were categorized as minimal for β less often than one third of the time, $BF_{10} = 111$. For all other proportions, the evidence was either ambiguous or in favor of the hypothesis that the true proportion was 1/3. Thus, based on our categorization method, we can conclude that our participants were relatively ideal for the kinds of information use related to β and P_{IG} , but were fairly uniform for α .

To address the concern that our results might be unduly influenced by the participants who were not strongly categorized by the model (i.e. participants with small winning BF_C), we repeated the preceding analysis using only categorizations for which the winning BF_C was greater than the median winning BF_C for that parameter. In the strongly-categorized sample,

we found quite similar results to the results from the whole sample. The main difference was that, for all parameters, the moderate categorization was reliably selected less often than 1/3 of the time, which was not the case for the full sample. This is because participants who were categorized as moderate tended to have smaller winning BF_C than the participants who were categorized as minimal or ideal, which, given our understanding of the categorization procedure, seems to be an empirical finding rather than an artifact of the categorization procedure. The main substantive outcome of this secondary analysis was that, when examining participants who were strongly categorized, α was not moderate, suggesting a bimodal distribution of minimal and ideal information users. However, as can be seen in Figure 8, the pattern of results is similar for both the full sample and the strongly-categorized participants.

Note that the fact that information use categorizations for α were uniformly distributed (at least in the full sample) does not imply that α is a meaningless parameter. On the contrary, α is a substantial correlate of WM capacity, attention, and informed guessing, as can be seen in Table 4, which suggests that it is a substantively important parameter. Finally, the Bayes factors that were used to categorize participants were stronger for α than for the other information use parameters, which suggests that the categorizations for α are, if anything, more meaningful than for the other parameters.

How Making Assumptions about Information Use Affects K Estimates

It is natural to think that making minimal or ideal assumptions about information use will have an effect on WM capacity estimates. For example, if participants are incorrectly assumed to be fully ideal information users, we would expect WM capacity to be underestimated. This is because, to achieve a certain level of performance on a task, one can either have more information or one can use that information more efficiently. This tradeoff means that if information use is assumed to be ideal, WM capacity estimates will decrease, and if information use is assumed to be minimal, WM capacity estimates will increase. We performed an analysis to measure how much the K estimates would change when it was assumed that information use was either completely minimal or completely ideal. We did this by changing α , β , and P_{IG} from free parameters to fixed values, either 0 for the completely minimal assumption or 1 for the completely ideal assumption, and then reestimating the remaining free parameters of the model. To save time, these analyses used 20,000 iterations of a Markov chain.

For the CD task, when minimal information use was assumed, the mean K estimate was 2.59, versus 1.57 when ideal information use was assumed, a difference of 1.02 objects. For the FM task, when minimal information use was assumed, the mean K estimate was 1.93, versus 1.41 when ideal information use was assumed, a difference of 0.52 objects. As can be clearly seen, there is a large difference in K estimates depending on whether a minimal or an ideal assumption about information use is made. In our primary analysis, when no assumptions were made about information use, the mean CD K estimate was 2.09 and the mean FM K estimate was 1.54. Thus, in the FM task, the K estimates under the assumption of ideal information use were closer to our best estimates for K than were the K estimates

under the minimal information use assumption. In the CD task, both the minimal and ideal assumptions were about the same distance, about half an object, from the best K estimates.

Discussion

In this study, we applied rigorous psychological process modeling techniques to learn about how minimally, moderately, or ideally participants use information in two WM tasks. We were able to estimate continuous measures of information use which give us a through understanding of how participants use information that they have in WM. Our results suggest a complex picture of information use in WM tasks. Based on our results, it is impossible to characterize all participants as either minimal, ideal, or moderate in terms of information use. This is in contrast to Chen and Cowan (2013), who had only less-well-developed methods and concluded that participants in a similar task were basically minimal information users. Instead, there seems to be a great deal of variability in how participants use information in our two tasks. However, one general trend was that participants were more ideal than minimal, at least for β and P_{IG} .

The most important implication of this study is our finding that participants cover a range of levels of information use is that models of WM cannot assume only ideal or minimal information use for all participants and expect the K estimates from those models to be accurate. One way to deal with this issue is to attempt to estimate how thoroughly participants are using information at test, as we have done here. By estimating information use it is possible to control for that variable when estimating K , even if information use is not of direct interest. In future research, if it is impossible or impractical to estimate information use in a study primarily interested in estimating K , our findings indicate that simply picking an assumption about information use could result in substantial error in the estimation of K . We would argue that if one had to pick between minimal and ideal information use assumptions for the FM task, assuming ideal information use would be the better choice. For the CD task, choosing to assume either minimal or ideal information use is a poor choice. A better choice for both tasks might be to use our estimates of the mean levels of information use, which are given in Table 3. Based in our sample, which was relatively large and heterogeneous for a study in cognitive psychology, we think that these information use estimates could be useful as norms when no other options are available.

Bounded Rationality

This study provides some support for the idea that humans have a bounded rationality that is constrained by information storage and processing limitations. In our tasks, we put participants under a memory load in a situation in which they would perform optimally if they were able to effectively reason about the contents of their memory. The kind of reasoning that participants were required to use to perform optimally on our tasks is not particularly demanding, yet according to the information use parameter estimates, our participants were not fully optimal in their reasoning. In addition, although overall our participants tended to be more ideal than minimal in information use, there was still a respectable proportion of our participants who were categorized as minimal information users. Our results relate to the idea that WM and reasoning share a common resource (e.g.

Baddeley & Hitch, 1974; Kane et al., 2004; Süß et al., 2002) and when WM is loaded with information, there is less of that common resource left for reasoning about the contents of WM. Similarly, Cowan, Morey, AuBuchon, Zwilling, and Gilchrist (2010) found that the process of allocating attention among stimuli based on task relevance was degraded in young children only when working memory was overloaded. It may be the case that full rationality can only be achieved by reducing memory load. Our results seem to be consistent with the idea of bounded rationality, in that our participants may have been rational to the best of their abilities given the memory and time constraints they were placed under. More work could be done to determine if bounded rationality describes behavior in our tasks and whether the amount of information in WM and how ideally that information is used can be shown to trade off with one another.

WM Capacity Estimates

The fact that the means of the K estimates for the two tasks were different (2.09 for the CD task and 1.54 for the FM task) might appear to pose a problem for the validity of our models, because the stimuli for both tasks were the same, so it might be expected that the K estimates should be the same. We believe that this is not a reason to doubt the validity of the models. The K parameters are estimated from participant responses which –although based on WM capacity – do not perfectly reflect the number of items that can be retained in working memory before a response is attempted. WM capacity estimates vary depending on stimulus and task characteristics. For example, Table 2 in Cowan (2001) lists capacity estimates from a number of studies that were selected to be quite similar to one another in terms of stimulus and task demands, yet the WM capacity estimates vary by more than 1 item between studies. The test situations in our tasks are substantially different from one another, so there is no reason to assume that the K estimates from the two tasks should be identical. Our K estimates from our two tasks are positively correlated and are of similar magnitude, so our models do not clearly disagree with one another.

In addition, we have an explanation for the difference in mean K estimates, which comes from response time data in our tasks. Mean response times in the FM task (2428 ms) were on about 1 second longer than mean response times in the CD task (1424 ms). This suggests that additional processing is needed in the FM task relative to the CD task, during which some information in WM could be lost due to decay (Ricker & Cowan, 2010) or interference (Oberauer, Lewandowsky, Farrell, Jarrold, & Greaves, 2012). It is well known, from studies on cognitive load, that memory maintenance is hurt by attention-demanding, concurrent tasks (e.g. Barrouillet, Bernardin, Portrat, Vergauwe, & Camos, 2007). Reasoning about the relationship between memory items and test probes is a kind of attentionally demanding task that must occur concurrently with the maintenance of the memory items that are being reasoned about, so we might expect that in both of our tasks that there would be some loss of memory items during the test phase. Because the FM task test phase takes longer to complete and includes more response objects that could cause interference, it could be that more information is lost for the FM task than for the CD task, which explains the discrepancy between the K estimates for the tasks.

The second author of this article has argued that items in WM are protected from decay or interference by being kept in the focus of attention (e.g. Cowan, 2001). However, more recent work has suggested that items in WM can be offloaded into the activated portion of long-term memory, where they might suffer decay or interference (Cowan, Sauls, & Blume, 2014). In addition, there is some evidence for both decay (Ricker, Spiegel, & Cowan, 2014) and interference from previous trials (Ricker et al., 2014; Shipstead & Engle, 2013) in visual array tasks such as ours. Thus, there is support for the idea that information in WM could be lost in our tasks. In addition, support for the idea that items in WM must be actively maintained comes from studies that used a post-cuing paradigm (Lepsien & Nobre, 2007; C. C. Morey, Morey, van der Reijden, & Holweg, 2013). It is possible that in our tasks, active maintenance cannot be performed while participants are reasoning about the test probes, which is more pronounced in the FM task. To conclude, our suggestion that the K estimates from our tasks may differ due to differential loss of memory items is supported by advances in theory resulting from recent studies.

Limitations of this Research

A limitation of the models used in this study is that the K parameters supposedly reflect only the number of objects for which both features are known. However, there are a few ways in which participants could have knowledge of only one feature of an object and still be able to use that knowledge to perform the task, which would most likely increase the K estimates to some extent. For example, a participant might know that in the sample array, the color red went with one object and the orientation 90° went with another object, but not know the other features of those objects. At test in the FM task, if the participant was probed with the color red, he or she could reject the response option 90° because he or she knows that those features did not go together, even though the participant does not know the other feature those features went with. Similarly, in the CD task test condition if the same information were in WM, if the probe object was red and oriented at 90° , the participant would know that there had been a recombination. We do not know how well participants are able to reason about objects with partial information, if they are able to at all, but there is good evidence that participants can have partial information about visual objects in WM (Fougnie & Alvarez, 2011; Hardman & Cowan, 2015; Oberauer & Eichenberger, 2013). To limit the kinds of within-feature partial information, though, as discussed below, we used clearly discriminable colors and orientations rather than continuous features.

Logically, the kind of reasoning required to make use of partial object information should have higher cognitive demands than reasoning about full objects. The reason is that full object reasoning only requires that participants analyze a single object in WM with respect to the test display, whereas in order to use information about partial objects, it is necessary to analyze two objects in WM at once with respect to the test display. This suggests that using partial object information may not occur with high frequency due to high cognitive demands. As the group-level model fits in Figure 5 show, the models are capable of fitting the observed data well with a reasonable number of parameters. We take this as evidence that the models are reasonable approximations of the psychological processes used by our participants and that the information use estimates we obtained are reliable. However, it is possible that our K estimates should not be interpreted as the number of objects in WM for

which both features are known, but as some combination of objects for which both features are known and objects for which only one feature is known.

The models we used were simplifications that do not capture the full nature of the way in which participants can perform these sorts of tasks. At the same time, all mathematical models are simplifications and they will be continue to be so until the models model millions of interconnected neurons. Regardless, we think that there is only limited potential for the un-modeled components of behavior to substantially affect our results. In order to estimate the number of objects for which different combinations of features is known, it will be necessary to use tasks that can measure knowledge of individual features as well as their conjunctions, which our tasks are unable to. Using only existing tasks and models, it is not easily possible to measure partial information about objects held in WM while also measuring simultaneous use of reasoning ability. Future research could attempt to estimate information use in tasks that allow for the separation of full- and partial-object knowledge with new tasks and models. For this study, based on the arguments we have made, it seems likely that the use of objects for which only partial information is known is uncommon and, as a result, it should not have substantially affected our findings.

Unexpected Results—There are a few results from this experiment that were unexpected. Primarily, it is unexpected that α and β do not correlate with one another given that, logically, they reflect very similar kinds of information use. Although we have provided some possible explanations for why we might not expect a correlation, the reasons why there should be a correlation are compelling. In follow-up research on this topic, this lack of correlation should be addressed, if possible. In addition, both β and P_{IG} were found to be fairly ideal, but this was not the case for α , which was found to be equally minimal and ideal. It is not clear why it would be the case that two kinds of information use are fairly ideal, but the third kind is not. Perhaps future research could examine whether participants could be made to be more ideal for α if they are trained to reason about the CD test situation. Another odd result is that u and P_{IG} from the CD task were negatively correlated, $r = -0.44$. The only explanation we have for this is that many participants had u that were higher than expected (u should be 0.5 if participants can accurately detect the rate at which recombined and unrecombined probes are used). Due to the negative correlation, higher u corresponds to lower P_{IG} . It could be that participants who were unable to accurately detect the proportion of probes of different types, resulting in high u , were also unable to effectively use informed guessing, resulting in lower P_{IG} . In spite of this explanation, we still think that this is an issue that could be explored further.

Precision of Feature Representations—One minor issue not addressed in our models is the potential for similar colors and orientations to be misremembered due to limited WM precision for features. Our models do not account for that possibility; instead, we chose to control for the issue by designing our stimuli to be highly distinguishable. To confirm that our stimuli are sufficiently distinguishable, we will use the results from a study that directly estimated the precision of WM for colors and orientations as norms. Fougny, Asplund, and Marois (2010) performed 4 experiments in which the precision of both color and orientation was estimated. The estimates were based on two set sizes, 3 and 6 items, and the units of

precision are in circular degrees as the colors were selected from a circular slice of a color space and the orientations were naturally circular in nature. Across the 4 experiments, the mean precision for color was 20.1 at set size 3 and 24.0 at set size 6 and the mean precision for orientation was 16.6 at set size 3 and 18.9 at set size 6. Our colors were chosen so as to be from an approximate circle in the HSB color space and the orientations can clearly be thought of as circular. Because we had 8 colors and 8 orientations, the distance in degrees between our stimuli is approximately 45 degrees for color and exactly 45 degrees for orientation. Given the norms from Fournie et al., our colors are about 2 standard deviations of WM precision from one another and our orientations are about 2.5 standard deviations from one another. As a result, errors resulting from imprecise WM representations should be rare in our study. Future studies that use our models should also use highly discriminable features so as to avoid the potential for imprecision to substantially affect the results of those studies.

Future Directions

A suggestion for future research would be to attempt to create a WM task that is well-suited to the analysis of how well information is used. This hypothetical task would ideally have two types of test with identical memoranda, where one test requires only memory and the other test requires both memory and reasoning about the items in memory. This would allow K to be estimated based on the memory-only condition and information use to be estimated based on the memory and reasoning condition. Having this sort of clean orthogonality between the parameters would help to improve the quality of the parameter estimates.

Our two tasks each have their own advantages and disadvantages. A disadvantage of the CD task is that the CD model must include a parameter related to guessing bias, u , but the FM task does not require this additional parameter. On the other hand, the FM task has only limited orthogonality between memory and reasoning because on any given trial we do not know whether participants were focused on searching for the correct response option or whether they were focused on eliminating incorrect response options. The CD task has the advantage over the FM task in that, on any given trial, we know whether the participant is making a decision about a recombined or unrecombined object. This additional knowledge possessed by the experimenter in the CD task about what sort of probe object the participant is making a decision about gives us a better understanding of what the participant is doing than we have in the FM task. However, the CD task still does not clearly separate memory and reasoning because participants themselves do not know whether they need to apply reasoning in an attempt to detect a recombination or whether they can simply focus on detecting unrecombined objects. It would benefit future research in this area if there was a WM task that better distinguished between memory and reasoning.

One modification of our tasks that could cause the CD and FM tasks to be more similar would be to change the FM task to only have two choices at test, i.e. a two-alternative forced-choice (2AFC) task. Using a 2AFC task would make it so that both the FM and CD tasks would always have two possible responses, regardless of the array size. In the FM task we used, the number of possible responses was always equal to the number of objects in the sample array, while there were only two possible responses in the CD task. In the present

experiments, only at array size 2 were the two tasks equated in terms of the number of possible responses that must be considered. It could be that using a 2AFC FM task would do a better job of equating the test situations in the CD and FM tasks in terms of complexity, causing response times and overall accuracy to be better equated at all array sizes. One outcome of this could be that the K estimates from the two tasks might be more comparable.

Final Conclusions

The results of this study do not support models of WM which assume that information use is either completely ideal or completely minimal. Instead, it appears that information use exists on a continuum from minimal to ideal, with different participants occupying different regions of the continuum. An important related result is that models that are used to estimate WM capacity can produce substantially inaccurate estimates of WM capacity if those models assume completely minimal or completely ideal information use. We suggest that researchers should strive to estimate how ideally participants use the information in WM in order to obtain the best possible estimates of WM capacity. It is possible that the concept of bounded rationality could help to explain why our participants were not fully ideal. However, at the same time, we found that our participants were generally more ideal than minimal. This suggests that even if our participants' information use was curtailed due to finite cognitive abilities, they were still able to regularly reason effectively about the contents of their WM. In sum, a lot can be gained by considering memory and reasoning in the same task and we have shown how it might be done for a large class of currently-used tasks.

Acknowledgments

Support was provided by NICHD Grant R01-HD21338 and the University of Missouri Life Sciences Fellowship Program.

Thanks go to Garrett Hinrichs and Suzanne Redington for help with data collection.

References

- Anderson JR. Is human cognition adaptive? *Behavioral and Brain Sciences*. 1991; 14:471–517.
- Baddeley, AD.; Hitch, G. Working memory. In: Bower, GH., editor. *The psychology of learning and motivation*. Vol. 8. New York: Academic Press; 1974. p. 47-89.
- Barrouillet P, Bernardin S, Portrat S, Vergauwe E, Camos V. Time and cognitive load in working memory. *Journal Of Experimental Psychology: Learning, Memory, And Cognition*. 2007; 33:570–85.10.1037/0278-7393.33.3.570
- Broadbent DA. Word-frequency effect and response bias. *Psychological Review*. 1967; 74:1–15. [PubMed: 5341440]
- Brown S, Heathcote A. Averaging learning curves across and within participants. *Behavior Research Methods, Instruments, & Computers*. 2003; 35:11–21.10.3758/BF03195493
- Chen Z, Cowan N. Working memory inefficiency: Minimal information is utilized in visual recognition tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2013; 39:1449–62.10.1037/a0031790
- Colom R, Abad FJ, Quiroga M, Shih PC, Flores-Mendoza C. Working memory and intelligence are highly related constructs, but why? *Intelligence*. 2008; 36:584–606.10.1016/j.intell.2008.01.002
- Cowan N. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*. 2001; 24:87–185. [PubMed: 11515286]

- Cowan N, Blume C, Saults S. Attention to attributes and objects in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2013; 39:731–747.10.1037/a0029687
- Cowan N, Hardman K, Saults JS, Blume CL, Clark KM, Sunday MA. Detection of the number of changes in a display in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. (in press).
- Cowan N, Morey CC, AuBuchon AM, Zwilling CE, Gilchrist AL. Seven-year-olds allocate attention like adults unless working memory is overloaded. *Developmental Science*. 2010; 13:120–133.10.1111/j.1467-7687.2009.00864.x [PubMed: 20121868]
- Cowan N, Saults JS, Blume CL. Central and peripheral components of working memory storage. *Journal of Experimental Psychology: General*. 2014; 143:1806–1836.10.1037/a0036814 [PubMed: 24867488]
- Danker JF, Gunn P, Anderson JR. A rational account of memory predicts left prefrontal activation during controlled retrieval. *Cerebral Cortex*. 2008; 18:2674–2685.10.1093/cercor/bhn027 [PubMed: 18321871]
- Engle RW, Tuholski SW, Laughlin JE, Conway ARA. Working memory, short-term memory, and general fluid intelligence: A latent-variable approach. *Journal of Experimental Psychology: General*. 1999; 128:309–331.10.1037/0096-3445.128.3.309 [PubMed: 10513398]
- Estes WK. The problem of inference from curves based on grouped data. *Psychological Bulletin*. 1956; 53:134–140. [PubMed: 13297917]
- Fougnie D, Alvarez GA. Object features fail independently in visual working memory: Evidence for a probabilistic feature-store model. *Journal of Vision*. 2011; 11:1–12.10.1167/11.12.3
- Fougnie D, Asplund CL, Marois R. What are the units of storage in visual working memory? *Journal of Vision*. 2010; 10:1–11.10.1167/10.12.27
- Gelfand, AE. Model determination using sampling-based method. In: Gilks, WR.; Richardson, S.; Spiegelhalter, DJ., editors. *Markov chain monte carlo in practice*. London: Chapman & Hall/CRC; 1996. p. 145–161.
- Gigerenzer G, Goldstein DG. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*. 1996; 103:650–669.10.1037/0033-295X.103.4.650 [PubMed: 8888650]
- Hardman K, Cowan N. Remembering complex objects in visual working memory: Do capacity limits restrict features of objects? *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2015; 41:325–47.10.1037/xlm0000031
- Hastings WK. Monte carlo sampling methods using markov chains and their applications. *Biometrika*. 1970; 57:97–109.10.1093/biomet/57.1.97
- Johnson JC. A test of the sophisticated guessing theory of word perception. *Cognitive Psychology*. 1978; 10:123–153.10.1016/0010-0285(78)90011-7 [PubMed: 668297]
- Kahneman D. A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*. 2003; 58:697–720.10.1037/0003-066X.58.9.697 [PubMed: 14584987]
- Kane MJ, Hambrick DZ, Tuholski SW, Wilhelm O, Payne TW, Engle RE. The generality of working memory capacity: A latent-variable approach to verbal and visuospatial memory span and reasoning. *Journal of Experimental Psychology: General*. 2004; 133:189–217.10.1037/0096-3445.133.2.189 [PubMed: 15149250]
- Lepsien J, Nobre AC. Attentional modulation of object representations in working memory. *Cerebral Cortex*. 2007; 17:2072–2083. [PubMed: 17099066]
- Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*. 1953; 21:1087–92.10.1063/1.1699114
- Miller GA. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*. 1956; 63:81–97.10.1037/h0043158 [PubMed: 13310704]
- Morey CC, Morey RD, van der Reijden M, Holweg M. Asymmetric cross-domain interference between two working memory tasks: Implications for models of working memory. *Journal of Memory and Language*. 2013; 69:324–348.10.1016/j.jml.2013.04.004
- Morey, RD.; Rouder, JN. *BayesFactor: Computation of bayes factors for common designs* [Computer software manual]. 2015. Retrieved from <http://CRAN.R-project.org/package=BayesFactor> (R package version 0.9.11-1)

- Nuijten, MB.; Wetzels, R.; Matzke, D.; Dolan, CV.; Wagenmakers, E-J. Bayesmed: Default bayesian hypothesis tests for correlation, partial correlation, and mediation [Computer software manual]. 2015. Retrieved from <http://CRAN.R-project.org/package=BayesMed> (R package version 1.0.1)
- Oberauer K, Eichenberger S. Visual working memory declines when more features must be remembered for each object. *Memory and Cognition*. 2013; 41:1212–27.10.3758/s13421-013-0333-6 [PubMed: 23716004]
- Oberauer K, Lewandowsky S, Farrell S, Jarrold C, Greaves M. Modeling working memory: An interference model of complex span. *Psychonomic Bulletin and Review*. 2012; 19:779–819.10.3758/s13423-012-0272-4 [PubMed: 22715024]
- Pashler H. Familiarity and visual change detection. *Perception & Psychophysics*. 1988; 44:369–378.10.3758/BF03210419 [PubMed: 3226885]
- Pratte MS, Rouder JN, Morey RD. Separating mnemonic process from participant and item effects in the assessment of ROC asymmetries. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2010; 36:224–232.10.1037/a0017682
- Prizmetel W, Lyon CA. The word-detection effect: Sophisticated guessing or perceptual enhancement? *Memory & Cognition*. 1996; 24:331–41.10.3758/BF03213297 [PubMed: 8718767]
- R Core Team. R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria: 2015. Retrieved from <http://www.R-project.org/>
- Ricker TJ, Cowan N. Loss of visual working memory within seconds: The combined use of refreshable and non-refreshable features. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2010; 36:1355–68.10.1037/a0020356
- Ricker TJ, Spiegel LR, Cowan N. Time-based loss in visual short-term memory is from trace decay, not temporal distinctiveness. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2014; 40:1510–1523.10.1037/xlm0000018
- 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; 105:5975–9.10.1073/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 and Review*. 2011; 18:324–30.10.3758/s13423-011-0055-3 [PubMed: 21331668]
- Shanks DR, Tunney RJ, McCarthy JD. A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*. 2002; 15:233–250.10.1002/bdm.413
- Shipstead Z, Engle RW. Interference within the focus of attention: Working memory tasks reflect more than temporary maintenance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2013; 39:277–289.10.1037/a0028467
- Suchow JW, Fougner D, Brady TF, Alvarez GA. Terms of the debate on the format and structure of visual memory. *Attention, Perception, & Psychophysics*. 2014; 76:2071–9.10.3758/s13414-014-0690-7
- Süß HM, Oberauer K, Wittman WW, Wilhelm O, Schulze R. Working memory capacity explains reasoning ability – and a little bit more. *Intelligence*. 2002; 30:261–88.10.1016/S0160-2896(01)00100-3
- Treisman A, Zhang W. Location and binding in visual working memory. *Memory & Cognition*. 2006; 34:1704–19.10.3758/BF03195932 [PubMed: 17489296]
- Tversky A, Kahneman D. Judgment under uncertainty: Heuristics and biases. *Science*. 1974; 185:1124–31.10.1126/science.185.4157.1124 [PubMed: 17835457]
- Tversky A, Kahneman D. The framing of decisions and the psychology of choice. *Science*. 1981; 211:453–58.10.1126/science.7455683 [PubMed: 7455683]
- Vulkan N. An economist's perspective on probability matching. *Journal of Economic Surveys*. 2000; 14:101–18.10.1111/1467-6419.00106
- Zhang H, Xuan Y, Fu X, Pylyshyn ZW. Do objects in working memory compete with objects in perception? *Visual Cognition*. 2010; 18:617–640.10.1080/13506280903211142

Appendix A

Informed guessing in the CD task

It is possible to separate guessing in a CD task into two kinds of guessing: informed guessing and uninformed guessing. Rouder et al. (2011) separated the two kinds of guessing for a different kind of CD task than we used in this study. Here we will provide an informed guessing formulation for the CD task we used in this study, in which the probe can be a recombination of features from multiple sample array items. For our CD task, uninformed guessing is the rate at which participants guess that there has been a recombination of features when they have no information about whether the probe object is or is not a recombination of features. An uninformed guess would happen when the participant has an attention lapse and fails to encode any information from the sample array. In Figures 3 and 4, an uninformed guess happens from the nodes labeled !A. Informed guessing requires that a participant use metamemory about the contents of their WM and reasoning abilities, along with their uninformed guessing rate, to calculate the informed guessing rate.

The basic idea is that even if participants do not have the information they need in order to definitively answer whether or not the probe object is a recombination, they may be able to guess more accurately if they consider what they do know and, consequently, may have been able to detect about the probe object. For example, let us assume that a participant studied 4 items in a sample array and successfully encoded 2 of them into WM. In addition, let us assume that when the participant was tested, the probe object was one of the two sample stimuli that the participant did not encode. In this case, the participant cannot detect the true state of the probe object regardless of whether it is or is not a recombination, so he or she must guess. Let us assume that the participant used information ideally on the trial. The participant knows that neither the color nor the orientation of the probe object match either of the objects that he or she has in WM. If the probe object was an unrecombined object, the participant knows that he or she would have recognized the probe object if it were either of the two objects in his or her WM. Thus, if the probe object were an unrecombined object, he or she would have detected it with probability 2/4. The participant also knows that only 2 objects were not known. Thus, if there had been a recombination, the participant would have only not detected it if the recombined features were from the two unknown objects. There

are $\binom{4}{2} = 6$ ways in which objects could be selected for a recombination, where $\binom{A}{B}$ is the choose function, choosing B objects from a set of A objects. There is only 1 way to select 2 objects out of which to create the probe object from the 2 unknown objects. As a result, the participant knows that had there been a recombination, he or she would have detected it with probability 5/6. Thus, the participant knows that he or she had a 1/2 chance of detecting an unrecombined object and a 5/6 chance of detecting a recombined object. Given that the participant did not detect either type of object, it seems more likely that an unrecombined object, rather than a recombined object, was missed, so the participant should be more likely to guess that the probe object was an unrecombined object.

We will assume that each participant knows the amount of information that they have in WM (i.e. K) and if they had used information ideally on that trial (i.e. if the participant

reached node I_1 in Figure 3 and attempted to detect recombinations). We also assume that the participant knows when an attention lapse has occurred, in which case the number of items in WM would be 0 for that trial. We also assume that a participant is aware of the total number of objects in the sample array, N , for the given trial. Finally, we assume that the participant has a subjective belief about the probability of the probe object being a recombined object, which is estimated by the model with the u parameter. Finally, we assume that participants probability match rather than choosing the most likely outcome when guessing.

To work out what the informed guessing rate should be, we set up an instance of Bayes' theorem for our particular case, just as was done by Rouder et al. (2011). We want to know the participant's subjective probability that there has been a recombination when the true state of the probe object has not been detected (i.e. the participant does not detect either a recombination or an unrecombined object). We will denote this $p(R|\bar{D})$, where R is the event that there has been a recombination and $\bar{D}$ is the event that the true state of the probe has not been detected (the line over D indicates logical negation of the event that the true state of the probe was detected). Conditioning on nothing being detected, $\bar{D}$, is just saying that we are only interested in cases in which the participant is guessing. Had the participant detected the true state of things, he or she would not be guessing and informed guessing would not apply. The full Bayes' theorem for this case is

$$p(R|\bar{D}) = \frac{p(\bar{D}|R)p(R)}{p(\bar{D})} \quad (5)$$

where $p(R|\bar{D})$ has been described above, $p(\bar{D}|R)$ is the probability that the true state of the probe has not been detected given that there has been a recombination, $p(R)$ is the subjective probability of a recombination, unconditional on detecting anything, and $p(\bar{D})$ is the subjective probability of not detecting the true state of the probe, regardless of whether a recombination has occurred. Because it is repetitive to continue to state that all probabilities are subjective, that fact will be assumed from here on.

Our goal is to express all four of the terms in this equation in terms of K , α , u , and N , all of which are either estimated or known. We are trying to solve for the informed guessing rate, which we will denote γ , which is the probability that a recombination has occurred given that the true state of the probe has not been detected, i.e. $p(R|\bar{D}) = \gamma$. The probability that a recombination has occurred, unconditional on detecting the true state of the probe, $p(R)$, is simply the uninformed guessing rate, u , so $p(R) = u$.

Turning to the third term in Equation 5, $p(\bar{D}|R)$ is the probability of not detecting the true state of the probe given that there has been a recombination. This is equivalent to $1 - p(D|R)$, in which $p(D|R)$ is the probability of detecting a recombination. Given the model presented in Figure 3, it is clear that $p(D|R)$ can be expressed as

$$p(D|R) = \alpha \left[\frac{K}{N} + \left(1 - \frac{K}{N}\right) \left(\frac{K}{N-1}\right) \right] \quad (6)$$

Note that throughout this appendix, it is assumed that $K \leq N$, so if $K > N$, then K must be constrained accordingly. $p(D|R)$ can alternatively be expressed with a ratio involving choose functions

$$p(D|R) = \alpha \left[\frac{\binom{N}{2} - \binom{N-K}{2}}{\binom{N}{2}} \right] \quad (7)$$

where $\binom{N}{2}$ is the number of the possible ways in which a recombination could have occurred and $\binom{N-K}{2}$ is the number of recombinations which would go undetected. Thinking of $p(D|R)$ in terms of choose functions can be helpful and we were able to put it to use in the example in the introduction of this appendix.

The fourth and final term in Equation 5 is the denominator, $p(\bar{D})$. By the law of total probability, $p(\bar{D})$ can be expressed as $p(\bar{D}|\bar{R})p(\bar{R}) + p(\bar{D}|R)p(R)$. We know that $p(R) = u$ and $p(\bar{R}) = 1 - p(R) = 1 - u$. In the previous paragraph, we worked out what $p(\bar{D}|R)$ is equal to. Finally, $p(\bar{D}|\bar{R})$ is the probability of not detecting that the probe object is an unrecombined object, which is simply $1 - K/N$ (see Figure 4 to see that this is true).

To summarize, we have the following equations that give all of the quantities in Equation 5 in terms of K , α , u , N , and, γ :

$$p(R|\bar{D}) = \gamma \quad (8)$$

$$p(R) = u \quad (9)$$

$$p(D|R) = \alpha \left[\frac{K}{N} + \left(1 - \frac{K}{N}\right) \left(\frac{K}{N-1}\right) \right] \quad (10)$$

$$p(\bar{D}|R) = 1 - p(D|R) \quad (11)$$

$$p(\bar{D}|\bar{R}) = 1 - \frac{K}{N} \quad (12)$$

$$p(\bar{D}) = p(\bar{D}|R)u + (1 - K/N)(1 - u). \quad (13)$$

We can substitute those quantities into Equation 5, which gives us

$$\gamma = \frac{p(\bar{D}|R)u}{p(\bar{D}|R)u + (1 - K/N)(1 - u)} \quad (14)$$

where $p(\bar{D}|R)$ is most convenient to include in its unexpanded form and can be calculated based on Equations 10 and 11. Implicit in these equations is the fact that K/N and $K/(N - 1)$ must be constrained to be in the interval $[0, 1]$. In addition, it should be understood that these equations should be applied to each array size to result in a different γ rate at each array size.

We can now give the equations for γ_M and γ_I . This can be done by treating α in Equation 10 as 0 if the participant is in a minimal guessing state or 1 if the participant is in an ideal guessing state. First, for the minimal guessing case

$$p(D|R, M) = 0 * \left[\frac{K}{N} + \left(1 - \frac{K}{N}\right) \left(\frac{K}{N-1}\right) \right] = 0 \quad (15)$$

$$p(\bar{D}|R, M) = 1 - 0 = 1 \quad (16)$$

$$\gamma_M = \frac{1 * u}{1 * u + (1 - K/N)(1 - u)} \quad (17)$$

where $p(D|R, M)$ comes from Equation 10 with α set to 0 (i.e. $p(D|R)$ given that the participant is in a minimal guessing state), and the two following equations are Equations 11 and 14 with the appropriate substitutions. The ideal case proceeds in the same way:

$$p(D|R, I) = 1 * \left[\frac{K}{N} + \left(1 - \frac{K}{N}\right) \left(\frac{K}{N-1}\right) \right] \quad (18)$$

$$p(\bar{D}|R, I) = 1 - p(D|R, I) \quad (19)$$

$$\gamma_I = \frac{p(\bar{D}|R, I)u}{p(\bar{D}|R, I)u + (1 - K/N)(1 - u)} \quad (20)$$

It should be noted that γ_M and γ_I are the perfect forms of the informed guessing that might be used by participants. Presumably, it takes some amount of effort to perform the mental calculations required to use informed guessing, so participants may not do so, at least not all the time. For this reason, we used the P_{IG} parameter to estimate the probability of informed guessing, versus uninformed guessing, that is used by participants. For completeness, we repeat here the equations from the main text that give the equations for the final guessing rates assumed to be used by participants once P_{IG} is factored in:

$$g_M = P_{IG} * \gamma_M + (1 - P_{IG})u \quad (21)$$

$$g_I = P_{IG} * \gamma_I + (1 - P_{IG})u. \quad (22)$$

These equations are straightforward mixtures of the perfect informed guessing rates, γ_M and γ_I , and the uninformed guessing rate, u . g_M and g_I should be calculated at each array size for each participant.

In our model, if the participant is not paying attention on a trial, he or she guesses that a recombination occurred with probability u . It is straightforward to show that this is not an *ad hoc* decision, but rather that it comes out of Equation 14. When the participant is not paying attention, the number of items in WM for that trial is 0. If that is the case, then $p(D|R)$, given in Equation 10, is also 0 and $p(\bar{D}|\bar{R})$ is 1. As a result, the equation for the informed guessing rate in the case of inattention, γ_I , is

$$\gamma_I = \frac{1 * u}{1 * u + (1 - 0/N)(1 - u)} = \frac{u}{u + (1 - u)} = \frac{u}{1} = u. \quad (23)$$

Next, we want to apply P_{IG} to γ_I to obtain the final guessing rate used by a participant in the case of inattention, g_I , just as we did for g_M and g_I . This is made quite simple because $\gamma_I = u$, allowing the following solution

$$g_I = P_{IG} * \gamma_I + (1 - P_{IG})u = P_{IG} * u + (1 - P_{IG})u = u. \quad (24)$$

In Figures 3 and 4 we used u instead of g_I for the sake of simplicity.

There is one important assumption that is made by this informed guessing formulation, which is that participants are accurately aware of K , α , and u , which is not necessarily the case. Because participants could easily be inaccurate about their WM abilities (c.f. Cowan et al., in press), it would not be surprising if participants did not use our informed guessing formulation perfectly. This could result in bias in the estimates of P_{IG} and potentially other parameters as well. However, imperfect use of informed guessing could simply result in low values for P_{IG} .

Appendix B Detailed Model Descriptions

We will treat responding on the CD and FM tasks as a binomial process, with a parameter that controls the probability of “success” on a trial, where a success is defined somewhat arbitrarily depending on the task. For the FM task, we define a success to be the event that a participant selected the correct response option. We will say that the observed number of correct responses for the i th participant at the j th array size, Y_{ij} , follows a binomial distribution

$$Y_{ij} \sim \text{Binomial}(p_{ij}, T_{ij}) \quad (25)$$

where p_{ij} is the probability that the i th participant makes a correct response at the j th array size and T_{ij} is the number of trials performed by the i th participant at j th array size. We have no direct interest in p , so we define it as a function of the parameters of interest in the model. For the FM task, the probability of making a correct response, p_{ij} is

$$p_{ij} = A_i \left\{ c_{ij} + (1 - c_{ij}) \left[\beta_i \frac{1}{N_j - \min(K_i^{(FM)}, N_j - 1)} + (1 - \beta_i) \frac{1}{N_j} \right] \right\} + (1 - A_i) \frac{1}{N_j} \quad (26)$$

where A_i is the attention parameter, $K_i^{(FM)}$ is the capacity estimate for the FM task, N_j is the j th array size, and c_{ij} is given by

$$c_{ij} = \min \left(\frac{K_i^{(FM)}}{N_j}, 1 \right). \quad (27)$$

As with proportion correct in the FM task, hits and false alarms in the CD task follow binomial distributions

$$H_{ij} \sim \text{Binomial}(h_{ij}, T_j^{(R)}) \quad (28)$$

$$F_{ij} \sim \text{Binomial}(f_{ij}, T_j^{(NR)}) \quad (29)$$

where H_{ij} and F_{ij} are the observed number of hits and false alarms, respectively, $T_j^{(R)}$ and $T_j^{(NR)}$ are the number of trials on which there was a recombination of features or there was no recombination, respectively, and h_{ij} and f_{ij} are the probability of a hit or a false alarm, respectively, as calculated from the model parameters with

$$h_{ij} = A_i \{ \alpha_i [d_{ij} + (1 - d_{ij}) (e_{ij} + (1 - e_{ij}) g_{ij})] + (1 - \alpha_i) g_{Mij} \} + (1 - A_i) u_i \quad (30)$$

$$f_{ij} = A_i \{ (1 - d_{ij}) [\alpha_i * g_{ij} + (1 - \alpha_i) g_{Mij}] \} + (1 - A_i) u_i \quad (31)$$

$$d_{ij} = \min \left(\frac{K_i^{(CD)}}{N_j}, 1 \right) \quad (32)$$

$$e_{ij} = \min \left(\frac{K_i^{(CD)}}{N_j - 1}, 1 \right) \quad (33)$$

where d_{ij} and e_{ij} are definitions used for convenience, g_{Mij} and g_{Iij} are the minimal and ideal guessing biases for participant i at array size j (note that the capital M and I subscripts are not indexing over anything), which are given in Equations 21 and 22 in Appendix A, $K_i^{(CD)}$ is the WM capacity parameter for the i th participant in the CD task, and α_i is the information use parameter for the i th participant.

Bayesian Parameter Estimation

To help provide constraint on the parameter estimates, hierarchical priors were placed on the K parameters from both tasks, A , and u . It is not easy to place hierarchical priors on A and u because they exist on the interval $[0, 1]$. Thus, A and u were transformed with the logit transformation so that the transformed parameters exist on the interval $(-\infty, \infty)$ so that hierarchical priors could be easily placed on the transformed parameters.

$$A^* = \text{logit}(A) \quad (34)$$

$$u^* = \text{logit}(u) \quad (35)$$

Similarly, the K parameters exist on the interval $[0, M]$, where M is the highest array size. We used $K^{*(CD)}$ and $K^{*(FM)}$ as latent, unbounded capacity parameters that were then transformed to be in the range $[0, M]$ as follows

$$T_K(x) = \max(\min(x, M), 0) \quad (36)$$

$$K^{(CD)} = T_K(K^{*(CD)}) \quad (37)$$

$$K^{(FM)} = T_K(K^{*(FM)}) \quad (38)$$

where T_K is the transformation function for the K parameters. It is reasonable to think that participants might have a latent capacity greater than the range of array sizes that were used, but we could not observe it due to the limited range of array sizes. Additionally, a negative latent capacity could arise due to a participant consistently responding in the reverse of how they should respond due to misunderstanding the instructions. This of course does not mean that the participant actually has a negative WM capacity, but rather that the data imply a negative WM capacity. Since we know that negative capacities are impossible, we convert negative latent capacities to zero.

The $K^{*(CD)}$, $K^{*(FM)}$, A^* , and u^* for each participant were given identical hierarchical priors, as given in the following template in which P is a placeholder that is replaced with each of the four parameters in turn.

$$P_i \sim \text{Normal}(\mu_P, \sigma_P^2) \quad (39)$$

$$\mu_P \sim \text{Normal}(0, 10^6) \quad (40)$$

$$\sigma_P^2 \sim \text{Inverse Gamma}(.1, .1) \quad (41)$$

Note that P is subscripted with an i , indicating that there is a P for each participant, but that μ_P and σ_P^2 are not, indicating that μ_P and σ_P^2 are shared by all participants (i.e. that the P_i all come from one population).

Hierarchical priors were not placed on α , β , or P_{IG} because we wanted constant priors on those parameters in order to use the categorization procedure as described in the main text. Each of α , β , and P_{IG} were given a uniform prior on the interval [0, 1].

The parameters were estimated using Bayesian techniques. None of the parameters of the models have known or obvious conjugate priors, so the Metropolis-Hastings algorithm (Hastings, 1970; Metropolis, Rosenbluth, Rosenbluth, Teller, & Teller, 1953) was used to sample from the posterior distributions of the parameters. The Metropolis-Hastings procedure was tuned to provide moderate acceptance rates for candidate parameter values: 95% of the acceptance rates were between 0.45 and 0.72. The posterior distributions of the parameters were sampled with 100,000 iterations of a Gibbs sampler. Convergence was verified by using visual inspection of several parameters' chains. No burn-in period was used as convergence was generally very quick. The Gibbs sampler was written in C++ and the output of the Gibbs sampler was further analyzed in R.

Figure 1.

Summary of the method of the CD and FM tasks. In each trial of both tasks, a sample array containing 2 to 5 to-be-remembered objects placed on the circumference of an invisible circle was presented. The tasks diverged only at test (on the right). The test for the CD task involved the presentation of a single object, which in this example was made by recombining a color and orientation from two different objects from the sample array. The other kind of test that could happen in the CD task used a whole object from the sample array (no recombination). One type of test for the FM task required the participant to pair the probe feature (outlined triangle indicating an orientation) with the color that orientation was paired with in the sample array (response number 3 in the example). The other FM test reversed the feature type of the probe and responses, so that a single color was shown and the participant was to select the orientation with which that color was paired in the sample array.

Figure 2.

Multinomial process tree for the FM model. Starting at node S, with probability $1 - A$ the participant is not paying attention on the trial, will arrive at node !A with no information in WM and must guess by selecting a response option at random, guessing correctly with probability $1/N$. If the participant is paying attention, he or she reaches node A and has K objects in WM (assuming that $K < N$). With probability K/N , the participant knows the object from which the probe feature was drawn and so knows the correct response option, and so from node W he or she responds correctly. If the participant does not have the probed object in mind, he or she reaches node G and must guess. Node G is marked with a diamond because it marks the point at which the model diverges for ideal versus minimal responding. From node G, with probability β the participant enters an ideal guessing state (node I), otherwise he or she enters a minimal guessing state (node M). In the ideal guessing state, the participant deductively eliminates K incorrect response options before guessing from the remaining response options. In the minimal guessing state, the participant guessing randomly from among all of the response options. The expression K/N is shorthand for $\min(K/N, 1)$. The expression $1/(N - K)$ is shorthand for $1/(N - \min(K, N - 1))$.

Figure 3.

Multinomial process tree for the CD model in the case where there has been a recombination of features. Starting at node *S*, with probability $1 - A$, the participant is not paying attention, reaches node *!A*, and guesses that there has been a recombination with probability u . If the participant is paying attention, he or she reaches node *A*, from which, with probability $1 - \alpha$, the participant enters a minimal responder state (node *M*) and is unable to detect recombinations. From node *M*, the participant guesses that there has been a recombination with probability g_M . If instead, with probability α , the participant enters an ideal responder state (node *I*₁), he or she uses information about both features of the probe object to respond, only guessing if no recombination is detected. The expressions K/N and $K/(N - 1)$ are shorthand for $\min(K/N, 1)$ and $\min(K/(N - 1), 1)$, respectively. Given that there has been a recombination of features, correctly responding that there has been a recombination results in a hit and incorrectly responding that there has been no recombination is a miss.

Figure 4.

Multinomial process tree for the CD model when there has been no recombination of features. Starting at node S, with probability $1 - A$, the participant is not paying attention, reaches node !A, and must guess. Because the participant has no information in WM, he or she must guess at the uninformed guessing rate, u . With probability A the participant is paying attention, has encoded K objects into WM, and reaches node A. With probability K/N , the participant has the probe object in WM and respond that the probe object is unrecombined, resulting in a correct rejection (CR). With probability $1 - K/N$, the participant does not have the probed object in mind and reaches node G, from which he or she must guess. Depending on whether the participant attempted to use information ideally (with probability α) or minimally (with probability $1 - \alpha$), he or she incorrectly guesses that there has been a recombination with probability g_I or g_M , respectively, resulting in a false alarm (FA). The expression K/N is shorthand for $\min(K/N, 1)$. Given that there has not been no recombination of features, correctly responding that there was no recombination of features is a correct rejection and incorrectly responding that there has been a recombination is a false alarm.

Figure 5.

Plots of the data and model predictions (“Pred.” in the legend) that are based on the primary, Bayesian modeling approach. A. CD task hit rate, B. CD task false alarm rate, and C. FM task proportion correct. Error bars are 95% confidence intervals.

Figure 6.

Plots of group-level K estimates based on closed-form solutions for K at each array size, with different assumptions about information use. The FM Ideal and FM Minimal models assume β of 1 and 0, respectively. The CD Ideal and CD Minimal models assume α of 1 and 0, respectively. The CD Intermediate model is the model used by Chen and Cowan (2013) as their benchmark model. All models assume $A = 1$ and the CD models assume $P_{IG} = 0$. The closed-form solutions do not converge across tasks as was the case for Chen and Cowan (2013).

Figure 7.

Histograms of individual participant mean parameter estimates. The attention parameter (panel A.) is shared by both the CD and FM tasks. Panels B and C show the parameters from the FM task and Panels D through G show the parameters from the CD task. The y-axes indicate the number of participants within each bin of x-values and thick vertical lines mark the grand mean value for each parameter. For the information use parameters, α , β , and P_{IG} , a value of 0 indicates minimal information use and a value of 1 indicates ideal information use.

Figure 8. Plots related to the Bayesian categorization of participants as minimal, moderate, or ideal. A. Histogram of posterior distribution for α of a single participant. The horizontal line shows the prior density and dashed vertical lines indicate the categories boundaries. This participant was categorized as ideal; see the text for further explanation. B and C. Proportions of participants put into each category when all participants were included (B) or when only those participants who received a strong categorization (top half of a median split) were included (C). The horizontal dotted line is at 1/3, the proportion of participants in each category under the null hypothesis. Black dots above the bars indicate that Bayes factor greater than 3 for that proportion being different from 1/3.

Table 1

Definitions of Terms and Symbols

Term or Symbol	Definition
Object	As an operational definition, we consider an object to be a combination of a color and an orientation.
Feature	One aspect of an object that determines the identity of that object. In this experiment, the relevant features are color and orientation.
CD	Change-Detection.
FM	Feature-Matching.
Probe Object	In the CD task, the single object presented at test that the participant must judge to be either a recombination of features from two different objects from the sample array or a whole object (unrecombined) from the sample array.
Probe Feature	In the FM task, the single feature presented at test for which the participant must select the matching feature from the response options. In Figure 1, the probe feature is the outlined triangle in the test display.
Response Options	In the FM task, the features presented at test from which the participant should select the feature originally paired with the probe feature. In Figure 1, the response options are the 4 colored circles.
A	The proportion of trials on which the participant is paying attention.
N	The number of objects in the sample array.
K	The asymptotic WM capacity for objects. For a single trial, the number of objects in WM is $\min(K, N)$ if the participant is paying attention, or 0 otherwise.
α	The proportion of trials on which a participant used information ideally to detect recombinations at test in the CD task.
β	The proportion of trials on which a participant used information ideally at test to reject incorrect response options in the FM task.
P_{IG}	The proportion of trials on which the participant used informed guessing perfectly in the CD task. See Appendix A for more information on informed guessing.
u, g_M, g_I	In the CD task, probabilities that the participant will guess that the probe object is a recombination, given that the participant is guessing and was not paying attention (u) or was using information minimally (g_M) or ideally (g_I).

Table 2

Means and Standard Deviations of Task Data at each Array Size

Measure	Array Size			
	2	3	4	5
CD Hit Rate	0.93 (0.07)	0.86 (0.1)	0.84 (0.13)	0.82 (0.12)
CD False Alarm Rate	0.07 (0.06)	0.24 (0.13)	0.38 (0.17)	0.45 (0.18)
CD Proportion Correct	0.93 (0.05)	0.81 (0.08)	0.73 (0.09)	0.68 (0.1)
FM Proportion Correct	0.91 (0.06)	0.74 (0.11)	0.54 (0.11)	0.41 (0.1)
CD Response Time	1239 (305)	1441 (383)	1493 (437)	1522 (474)
FM Response Time	1759 (393)	2386 (585)	2696 (762)	2872 (963)

Values in table are presented as “mean (standard deviation)”. Response times are in ms.

Table 3

Means and Standard Deviations of the Primary Parameters

Task	Parameter	Mean (SD) [*]
Both	A	0.89 (0.03)
FM	K	1.54 (0.33)
FM	β	0.58 (0.13)
CD	K	2.09 (0.22)
CD	a	0.51 (0.21)
CD	P_{IG}	0.56 (0.14)
CD	u	0.70 (0.13)

^{*}: Mean (standard deviation) of the posterior means of the parameters.

Table 4

Between-Participant Parameter Correlations

	FM <i>K</i>	CD <i>K</i>	α	β	P_{IG}	<i>A</i>	<i>u</i>
FM <i>K</i>	1						
CD <i>K</i>	0.44 [*]	1					
α	0.55 [*]	0.75 [*]	1				
β	0.33	-0.04 [†]	0.10 [†]	1			
P_{IG}	0.16 [†]	0.29	0.45 [*]	0.08 [†]	1		
<i>A</i>	0.54 [*]	0.55 [*]	0.54 [*]	0.40 [*]	0.37 [*]	1	
<i>u</i>	0.04 [†]	-0.19 [†]	0.02 [†]	0.11 [†]	-0.44 [*]	-0.03 [†]	1

Each cell gives the correlation between the row and column parameters. For each parameter, the values correlated are the means of the posterior distributions for each participant. The Bayes factor for the hypothesis that the correlation was not 0 was greater than 3

^{*} ; evidence for a correlation, less than 1/3

[†] ; evidence against a correlation, or between 3 and 1/3 (no mark; indecisive evidence).