RUNNING HEAD: Bias Effects

Bias effects in a two-stage recognition paradigm: A challenge for 'pure' threshold and signal
detection models

Qiuli Ma, Jeffrey J. Starns
University of Massachusetts Amherst

David Kellen
Syracuse University

Correspondence:

Jeffrey Starns
Department of Psychological and Brain Sciences
University of Massachusetts
135 Hicks Way
Amherst, MA 01003-9271
jstarns@psych.umass.edu
PHONE: 413-545-5951

Abstract

We explored a two-stage recognition memory paradigm in which people first make single-item “studied”/“not studied” decisions and then have a chance to correct their errors in forced-choice trials. Each forced-choice trial included one studied word (“target”) and one non-studied word (“lure”) that received the same previous single-item response. For example, a “studied”-“studied” trial would have a target that was correctly called “studied” and a lure that was incorrectly called “studied.” The two-high-threshold (2HT) model and the unequal-variance signal detection (UVSD) model predict opposite effects of biasing the initial single-item responses on subsequent forced-choice accuracy. Results from two experiments showed that the bias effect is actually near zero and well out of the range of effects predicted by either model. Follow-up analyses showed that the model failures were not a function of experiment artifacts like changing memory states between the two types of recognition trials. Follow-up analyses also showed that the dual process signal detection (DPSD) model made better predictions for the forced-choice data than 2HT and UVSD models.

Bias effects in a two-stage recognition paradigm: A challenge for ‘pure’ threshold and signal detection models

A popular way to test memory ability is to ask someone to discriminate stimuli that were encountered in a target context from those that were not, such as deciding whether or not a certain word was studied in a previous list. This procedure, known as a recognition task, has been influential in both applied research (e.g., Prull, Dawes, Martin, Rosenberg, & Light, 2006) and theoretical research (e.g., Ratcliff, Sheu, & Gronlund, 1992). Researchers rely on decision models to link recognition performance to psychological processes such as memory acuity – the ability to discriminate studied and nonstudied items – and response bias – the overall predilection for responding “studied” versus “not studied” (Banks, 1970; Lockhart & Murdock, 1970). Models that make different assumptions about memory and decision processes can lead to different research conclusions (Rotello, Masson, & Verde, 2008), so it is critical for researchers to test the basic assumptions of decision models.

The original goal of the present work was to test two popular decision models – the two-high-threshold (2HT) model and the unequal-variance signal detection (UVSD) model – by challenging them to predict the effect of bias in a two-stage procedure for testing recognition memory (Starns, Dube, & Frelinger, 2018; Voormann, Rothe-Wulf, Starns, & Klauer, 2021). The two models predict bias effects that go in *opposite directions*, predictions that are not dependent on specific values for the model parameters. These predictions were tested with two experiments, the second of which was preregistered on the Open Science Framework. Critically, the models were not allowed to adjust any parameter values in an effort to accommodate the result that was used to test the models, so in contrast to common usage, the model predictions were *actually* predictions (Busemeyer & Wang, 2000). Overall, results showed that bias effects were near zero

and well out of the range of predictions of either model, providing evidence that both models fail to closely approximate the true decision process in recognition memory. This “plague on both houses” led us to broaden the set of candidate models. As it turns out, the observed bias effects are at odds with other models as well, even when some of them are allowed relaxed assumptions. Ultimately, the results were most consistent with a dual-process account in which responses to studied items result from a mixture of “familiarity” and “recollection” processes.

Model Overview

Researchers have explored many mathematical models of the decision process in recognition memory, and some of these models have been applied in thousands of studies over many decades (e.g., Egan, 1958; Green, & Swets, 1966; Swets, 1961; Snodgrass, & Corwin, 1988; Luce, 1963). One fundamental question in the modeling literature concerns the nature of the information retrieved from memory. There are two influential modeling approaches that starkly disagree on this question: multinomial processing tree (MPT) models and signal detection theory (SDT) models. MPT models hold that recognition decisions are informed by several discrete inner cognitive states (e.g., Batchelder & Riefer, 1999; Kellen, Erdfelder, Malmberg, Dube, & Criss, 2016; Luce, 1963; Swets, 1961; Snodgrass, & Corwin, 1988), whereas SDT holds that those decisions are the result of comparisons between decision criteria and memory strength values drawn from a continuous distribution (Egan, 1958; Green, & Swets, 1966).

For recognition memory, the most popular member of the MPT model class has been the two high-threshold (2HT) model (Bröder & Schütz, 2009; Bröder, Kellen, Schütz, & Rohrmeier, 2013; Snodgrass, & Corwin, 1988). According to this model, which is illustrated in Figure 1A, a studied word (alternatively, an “old” word or a “target”) can lead to two mental states. With probability D_o , evidence identifying the word as one from the study list is remembered, and the

Detect-Old mental state is reached. This leads to an “old”/“studied” response. With complementary probability $1-D_o$, no evidence is remembered and the *Guess* state is reached instead. In the case of a non-studied word (alternatively, a “new” word or a “lure”), the presence of disqualifying information (see Gallo, 2006, Chap. 5; Strack & Bless, 1994) leads to the *Detect-New* mental state being reached, which in turn results in a “new”/“not studied” response. If no disqualifying information is available, then the *Guess* state is reached. Once target or lure test words enter into the *Guess* state, they are treated equally; that is, the probabilities of guessing “old”/“studied” or “new”/“not studied”, which are denoted by g and $1-g$, respectively, are identical for both types of items. The critical property that defines the model as a “high”-threshold process is that targets never enter into the *Detect-New* state, and lures never enter into the *Detect-Old* state (for critical discussions of this assumption, see Kellen, Winiger, Dunn, & Singmann, 2021; Moran, 2016; Starns, 2021).

In contrast to the 2HT model, the Gaussian Signal-Detection model assumes a continuous distribution of memory strength for both targets and lures (Banks, 1970; Green, & Swets, 1966; Kellen & Klauer, 2018; Lockhart & Murdock, 1970; Macmillan, & Creelman, 2005; Wixted, 2007). As shown in Figure 1B, the mean of target distribution (dashed curve) is greater than that of the lure distribution (solid curve), reflecting the fact that memory evidence tends to be stronger for these items. A decision criterion, κ , is set along the dimension of memory strength to determine the recognition response: if the memory strength for an item falls on the right of the criterion, an “old” response is made, and if it falls on the left, a “new” response is made (this is the vertical dotted line in Figure 1B). A basic version of the SDT recognition memory model characterizes the memory strengths of targets and lures in terms of two Gaussian distributions with equal variances. In practice, however, it well known that an *Unequal Variance Signal*

Detection (UVSD) model, in which the variance of the distribution of target strengths is greater than its lure counterpart, is known to provide a better fit of recognition data (e.g., Egan, 1958; Wixted, 2007). Figure 1A displays the unequal-variance case. The means of the target and lure distribution are denoted by μ_t and μ_l , and their standard deviations by σ_t and σ_l . By convention, and without any loss of generality, the mean and standard deviation of the lure distribution are set to 0 and 1, respectively (Macmillan & Creelman, 2005). The criterion κ can move along the dimension ranging from the most liberal at the left end to the most conservative at the right end.

The 2HT and UVSD models have been pitted against each other in a number of different studies (e.g., Bröder & Schütz, 2009, Bröder et al., 2013; Chechile, 2012; Chen, Starns, & Rotello, 2015; Dube & Rotello, 2012, Dube, Starns, Rotello, & Ratcliff, 2012; Kellen & Klauer, 2014, 2015; Kellen, Klauer, & Bröder, 2013; Klauer & Kellen, 2015; Province & Rouder, 2012; Starns, 2020, Starns et al., 2018, Starns & Ma, 2018; Voorman et al., 2021). But whereas earlier comparisons mostly focused on goodness-of-fit indices (corrected for model flexibility) associated with Receiver-Operating Characteristic (ROC) functions, more recent work has targeted specific qualitative differences between the models (for a discussion on “critical testing”, see Kellen et al., 2021). The present work continues this recent trend. What is novel here is the way in which *responses biases*, which have mostly played a merely instrumental role by providing the means to construct ROC functions, take center stage. As we will show below, the 2HT and UVSD models make *opposing* predictions across manipulations of response bias. The details surrounding these predictions are reported in preregistration documents posted on the Open Science Framework (OSF) site before data collection began for Experiment 2 (review link: https://osf.io/kn2tw/?view_only=7c207f36faf745fe85ff47acc75ed648).

The remainder of this manuscript is organized as follows: First, we will describe the experimental paradigm used, and how the 2HT and UVSD models make opposing predictions. We will then discuss the results obtained from two experiments, and how both models fail to accommodate them. Finally, we will discuss a number of alternative models and their associated predictions.

The “Error Correction” Paradigm

The test developed here is based on a recently developed paradigm in which participants are given the opportunity to correct errors from previous single-item recognition trials in subsequent forced-choice trials (Starns et al., 2018; Voormann et al., 2021). After learning a list of words, participants completed a series of test blocks as summarized in Figure 2. Each block began with single-item recognition trials; that is, participants saw a single word appear on the computer screen and decided whether it was “studied” or “not studied.” We used payoffs to manipulate response biases on the single-item trials by establishing a “safe” and a “risky” response, and each block began with a screen identifying which response was safe and which was risky. The safe and risky responses appear in white and black in Figure 2, respectively (the actual colors for participants were green and red). When participants selected the safe response, they earned 3 points for correct responses and lost 1 point for errors. When they selected the risky response, they earned 1 point for correct responses and lost 3 points for errors. Participants were informed of this asymmetry, and they were advised to avoid making a risky response when they were unsure of their ability to answer correctly to avoid the extreme cost of risky errors. In the conservative condition, the “studied” response was risky and the “not studied” response was safe. In the liberal condition, the “not studied” response was risky and the “studied” response

was safe. Thus, we expected that participants would make more “old” responses overall in the liberal than the conservative condition, similar to previous recognition studies that used response payoffs to manipulate recognition biases (Curran et al., 2007; Healy & Kubovy, 1978; Henriques, Glowacki, and Davidson, 1994; Snodgrass & Corwin, 1988; Van Zandt, 2000).

Each block had 12 single-item trials that were followed by forced-choice trials. For the latter, participants were shown a target and a lure from the previous single-item recognition trials, one of which was previously classified incorrectly (Starns et al., 2018). Table 1 shows an example block of single-item and forced-choice trials to demonstrate how trials were constructed. Some forced-choice trials were “studied”-“studied” (S-S) trials, comprising a target that was correctly called “studied” and a lure that was incorrectly called “studied.” For the example in Table 1, “restaurant” was a lure with an incorrect single-item response, and it was paired with “way,” a word selected at random from the pool of targets that were correctly recognized, to make an S-S forced-choice trial. Some trials were “not studied”-“not studied” (N-N) trials, comprising a target that was incorrectly called “not studied” and a lure that was correctly called “not studied.” For Table 1, two words in the block – “spectacular” and “disaster” – were targets with incorrect single-item responses, and each was paired with a word randomly selected from the pool of lures that were correctly rejected to produce N-N trials. Note that the condition labels (S-S and N-N) refer to the previous *response* that the participant made for both items, not the items themselves – both S-S and N-N trials always included one target and one lure. In this forced-choice phase, participants were asked to choose the word that was actually on the study list, and were advised that every trial included one item that they previously misclassified.

Bias Effects Predicted by the 2HT Model

According to the 2HT model, performance on the forced-choice trials is determined by the relative proportion of detection- and guessing-based responses in the single-item recognition trials, which is in turn affected by the bias to respond “old/studied” or “new/not studied” from the *Guess* state. The 2HT model states that the probability that a target is judged “old” in single-item recognition is $D_o + (1 - D_o)g$; that is, the probability that a target will be detected as “old” plus the probability that detection will fail but the participant will guess “old”. Therefore, the proportion detection-based responses for targets called “old” is $P(D_o | \text{"old"}) = \frac{D_o}{D_o + (1 - D_o)g}$.

Notice that g is in the denominator, so $p(D_o | \text{"old"})$ decreases as g increases. What this means is that when participants are reluctant to guess “old,” the relative proportion of detection-based “old” responses is quite high. This proportion will decrease the more people are willing to guess “old”. The opposite pattern holds in the case of “new” responses for lures: The probability that lure is judged “new” is $D_n + (1 - D_n)(1 - g)$, and the proportion of detection-based responses among lures called “new” is $p(D_n | \text{"new"}) = \frac{D_n}{D_n + (1 - D_n)(1 - g)}$. Note the term $(1 - g)$ in the denominator, which establishes that the relative proportion of detection-based “new” responses among lures will be greater the more participants are reluctant to guess “new”.

These predictions have clear implications in terms of the predicted accuracy of the subsequent S-S and N-N trials.¹ Let us first consider the S-S trials: When presented with a target and a lure that were judged “old”, the probability of selecting the target is a function of the memory state associated with each item. If the target is in a Detect-Old state, then a correct

¹ We assume that the information available from memory for an item is the same when it is tested in the single-item trials and when it is tested in the forced-choice trials. In the experiments, there was a short lag between two test trials with the same item, so it is unlikely that the memory state changed drastically. We will return to this assumption in the follow-up analyses to show that the model predictions are robust to plausible changes in memory state from the single-item to forced-choice tests.

response (i.e., selecting the target) is produced with probability 1. But if the target is in the Guess state, then the probability of a correct response is .50. Please note that according to the 2HT, “old” responses for lures can only take place via guessing. The probability of a correct response in an S-S trial corresponds to:

$$P(\text{correct}, S - S) = \frac{D_o + .50(1 - D_o)g}{D_o + (1 - D_o)g}$$

The relationship between the probability of a correct response in S-S trials and parameters D_o and g is illustrated in the top-left panel of Figure 3, with different lines for different D_o values and variation in g on the x -axis. The probability of correct responses increases as g decreases, regardless of D_o (note that the x -axis is reverse coded).

Let us now turn to the case of N-N trials. A correct response will occur with probability 1 if the lure is in the Detect-New state, and with probability .50 if it is in the Guess state. The relationship between the probability of a correct response and parameters D_n and g is given by:

$$P(\text{correct}, N - N) = \frac{D_n + .50(1 - D_n)(1 - g)}{D_n + (1 - D_n)(1 - g)}$$

This relationship is illustrated in the middle-left panel of Figure 3. It is a “mirror image” of the relationship in S-S trials, given that they only differ in terms of how guessing affects the numerators and denominators. Finally, the relationships found in both S-S and N-N trials can be taken together in terms of a “difference of differences” (DoD; i.e., an interaction). The bottom-left panel of Figure 3 illustrates how the difference $P(\text{correct}, S-S) - P(\text{correct}, N-N)$ is affected when considering two different guessing probabilities. So, for example, a value of -0.20 on the x -axis means that the probability of guessing old was .20 lower in the conservative condition than the liberal condition. The different lines correspond to different combinations of D_o and D_n

values. In the case of the 2HT, *this difference of differences is always positive*, and takes on more extreme values as the response-bias effect increases.

Bias Effects Predicted by the UVSD Model

In the UVSD model, the probability of a correct responses in forced-choice trials is determined by the probability that memory strength is higher for the target item than the lure, where the distributions of target and lure strength are conditioned on their respective single-item responses. Thus, bias on the single-item trials can affect forced-choice accuracy by changing these conditional distributions. Figure 4 demonstrates this mechanism for S-S trials, with full strength distributions across all items in Panels A-C and the corresponding conditional strength distributions for recognized items in Panels D-F. The recognition response criterion becomes more liberal from Panels A-C and D-F. For S-S trials, both the target and the lure were called “old” in the single-item trials, which means that they both have strength values above the identification criterion (κ). So the distributions of memory strength for items that appear in S-S forced-choice trials are different than the original strength distributions, and they are shifted to higher strength values as shown. This selection effect increases the overlap in the target and lure strengths, especially when single-item responding is conservative (high κ). For example, in Panels A and D not many lures have strength values high enough the pass the conservative response criterion, but the ones that do are necessarily quite strong. Therefore, the conditional distributions overlap to a great deal, and forced-choice accuracy is relatively low (.72). As responding becomes more liberal, the selection effect subsides and the conditional strength distributions pull apart. Accordingly, forced-choice accuracy increases to .76 in Panels B and E and .80 in Panels C and F.

For N-N trials the selection effect works in the other direction, as all items in these trials were called “not studied” and thus have strength values below the single-item response criterion κ . Figure 5 shows conditional distributions for items called “not studied” with the same overall strength distributions and response criteria as Figure 4. The distributions that are conditioned on a “not studied” response (Panels D-F) tend to have lower strength values than the distributions across all items, of course, as all of these are items that fell below the recognition criterion. The conditionalization increases the overlap in the target and lure distributions: although targets are stronger than lures in general, selecting only items that are too weak to pass the recognition criterion degrades this advantage. As the response criterion becomes more liberal in Panels A-C, the selection effect grows more severe in Panels D-F and forced-choice accuracy decreases from .75 to .60.

The probability of responding correctly in an S-S trial is given by

$$P(\text{correct}, S - S) = \int_{\kappa}^{\infty} \frac{f(x, \mu_t, \sigma_t)}{1 - F(\kappa, \mu_t, \sigma_t)} \frac{F(x) - F(\kappa)}{1 - F(\kappa)} dx$$

In this equation, $f()$ indicates the Gaussian probability density function and $F()$ indicates the Gaussian cumulative distribution function. For example, $f(x, \mu_t, \sigma_t)$ is the probability density of strength value x on the target distribution given its mean μ_t and standard deviation σ_t . To get the probability density for *recognized* targets, one must normalize strength values by dividing by the total proportion of targets that are recognized; i.e., $1 - F(\kappa, \mu_t, \sigma_t)$. Thus, the first fraction in the equation gives the probability density at strength value x for a *recognized* target for any x above the recognition criterion (the probability density is zero for all values below the criterion, as these items have been excluded from the class of recognized targets). The second fraction is the probability that a *falsely recognized* lure has a strength value below x , and it is found by dividing the proportion of all lures between the criterion and x by the proportion of all lures above the

criterion. In other words, $F(x) - F(\kappa)$ is the probability that a lure has a strength value below x and above the recognition criterion, and $1 - F(\kappa)$ is the probability that a lure has a strength value above the recognition criterion. Dividing the two gives the probability that a lure *that was judged “old”* has a strength value below x . Multiplying the two fractions gives the joint probability density that a recognized target has strength value x and has a higher strength value than a randomly selected lure that was falsely recognized. Integrating over all x values above the recognition criterion gives the total probability that any recognized target would have a higher strength value than a recognized lure, corresponding to the predicted forced-choice accuracy.

Similarly, the probability of responding correctly in an N-N trial is given by:

$$P(\text{correct}, N - N) = \int_{-\infty}^{\kappa} \frac{f(x, \mu_t, \sigma_t)}{F(\kappa, \mu_t, \sigma_t)} \frac{F(x)}{F(\kappa)} dx$$

The first fraction takes the probability density for value x on the target strength distribution, $f(x, \mu_t, \sigma_t)$, and divides it by the proportion of the target distribution that falls below the recognition criterion, $F(x, \mu_t, \sigma_t)$. This yields the probability density for targets judged “new” (with densities for all strength values greater than κ equal to zero). The second fraction divides the proportion of the lure distribution below strength value x by the proportion below the response criterion, κ , yielding the conditional probability that a lure strength is less than x given that the lure was called “not studied.” Multiplying the two fractions gives the joint probability density that an unrecognized target has strength x and the unrecognized lure paired with it has strength below x . Integrating across all strength values below the response criterion gives the overall probability of correctly selecting the target in an N-N trial.

The right panels of Figure 3 illustrate the predictions associated with the UVSD model. First, we see that accuracy decreases in S-S trials (top-right) as the criterion becomes increasingly conservative. The opposite pattern is found in the N-N trials (middle-right), in

which accuracy is greater under more conservative criteria. What is happening here is that “bunching up” the strength distributions by conditioning on a particular response obscures the average strength difference between targets and lures, resulting in decreased accuracy in one type of forced-choice trial, along with an increase in other. Finally, the bottom-right panel shows predictions for the difference of differences (DoD). The x -axis is the criterion difference between bias conditions; for example, a value of 1 means that the conservative criterion is 1 above the liberal criterion, where the units are the standard deviation of the lure strength distribution. The different lines are different combinations of μ_t , μ_l , and σ_t values. The UVSD model generally predicts negative DoDs, which become more extreme as the response-bias effect increases.²

Interim Discussion

Previous work comparing the 2HT and UVSD has very often relied on “nested” predictions, with the UVSD expecting an effect in specific direction, whereas the 2HT model expects no effect whatsoever (i.e., a null effect; e.g., Chen et al., 2015; Kellen & Klauer, 2014, 2015; Kellen et al., 2021). In the present case, we see that both models make opposing (i.e., non-nested) predictions. Their distinct predictions stem from the way in which response bias is captured by each model: Whereas the 2HT model associates response bias with the guessing probability, g , the UVSD model attributes it to the criterion, κ , which determines how much strength is necessary in order to recognize a test item as “old”. Moreover, note how the predictions at the level force-choice accuracy “emerges” from the single-item judgments that

² Technically, it is possible for the UVSD model to predict positive DoDs. But it can only do so when positing “pathological” characterizations. More specifically, the prediction of positive DoDs emerges when assuming that the mean of the target strength distribution *is below* the mean of the lure strength distribution or that the two distributions have similar means with much greater variability in target strength than lure strength. Neither assumption is consistent with any known learning mechanism, nor is it observed when fitting the model to data at large.

precede them. This introduces another layer of constraint in model testing, in the sense that a successful model should be able to accurately predict effects at the level of forced-choice judgments based on their characterization of single-item judgments.

We tested model predictions with two experiments, the second of which was preregistered. The project's OSF page (review link: https://osf.io/kn2tw/?view_only=7c207f36faf745fe85ff47acc75ed648) has documents detailing our predictions, the planned methods, and the analysis strategy. The two experiments were near replications, with slight variation in the length of study and test lists. Accordingly, we will report combined results across both experiments to best represent the conclusions supported by all available evidence.

Method

Participants

We collected data from a total of 196 college students. Experiment 1 included 51 participants from the University of Massachusetts (UMass) Amherst and Experiment 2 included 50 participants from UMass Amherst and 95 participants from Syracuse University. We planned to have two sessions of data collection for each participant, but not all participants returned for Session 2. For Experiment 1, 43 participants completed both sessions and 8 completed only a single session. For Experiment 2, 123 participants completed both sessions and 22 completed one session. We kept data from single-session participants, because the two sessions were replications of the same task and our analyses relied on hierarchical Bayesian models that

automatically adjust for different trial counts. Participants were compensated with extra credit in their psychology courses.

As outlined in the preregistration document (review link: https://osf.io/kn2tw/?view_only=7c207f36faf745fe85ff47acc75ed648), participants were excluded from data analyses if the difference between the single-item hit rate and false-alarm rate was less than .10 and/or if they made the safe response less than half of the time in their initial guesses (indicating that they did not understand the bias instructions). Reviewers were (justifiably) concerned that our task might encourage guessing in the single-item recognition task. We analyzed RT data to assess that possibility and additionally excluded participants who made over 10% of their single-item recognition responses in less than 550 ms (see the *Results* section for the rationale behind this criterion). The decision to make these additional exclusions did not change any research conclusions.

For Experiment 1, 4 participants were excluded by the original criteria and no participants were excluded by the guessing criterion, leaving 47 participants eligible for analyses. For Experiment 2, 36 participants were excluded by the original criteria and an additional 13 participants were excluded by the guessing criterion, leaving 100 participants eligible for analyses. This exactly matches the intended sample size for Experiment 2 reported in the pre-registration document. Across both experiments, we had 147 participants eligible for analysis.

Materials

Words for the study and test lists were randomly sampled from 1098 nouns from the SUBTLEXus dataset (Brysbaert & New, 2009). The word frequency ranged from 10.02 to 99.49 per million words based on subtitles from American films and television series. Each study list had eight filler words, four at the beginning and four at the end, with 64 and 60 standard words

in between for Experiments 1 and 2, respectively. Each test started with 4 filler words, two targets from the early fillers on the study list and two filler lures not seen before, followed by standard words. Experiment 1 had 96 standard words on the test and Experiment 2 had 120, with each evenly split between targets and lures.³ There was no distinction between filler items and standard items from the participants' perspective, but the former were excluded from all analyses. The purpose of the filler items was to give participants a few trials to acclimate to each task and to ensure that no words from the end of the study list showed up in the first few test trials while they were still in working memory. We also created a practice cycle with 32 words on the study list and 52 words on the test list.

Procedure

The first session comprised one practice study-test cycle and two standard cycles. Thus, participants who completed only one session had 192 total single-item trials for Experiment 1 and 240 for Experiment 2. The second session was identical to the first, except that the first cycle was converted from a practice to a standard cycle. Thus, participants who returned for the second session completed five study-test cycles in all, producing 480 total single-item trials for Experiment 1 and 600 for Experiment 2. The number of forced-choice trials varied across participants based on their single-item performance as described below. Words were randomly assigned to each session and study-test cycle, and participants were (correctly) informed that words from a given cycle would never reappear in the experiment and that they would never have to remember words studied before the most recent study list.

³ We did not use all of the standard studied words as targets on the test list in Experiment 1 because of a coding vestige. The code was adapted from earlier experiments in which participants completed a recall task at encoding, and the words prompted for recall were excluded from the recognition test. Excluding these trials does not alter the task in any meaningful way, but it reduces the number of trials available for analysis. For Experiment 2, we changed the code to include all standard words from the study list as target items on the test.

During the study lists, participants were asked to remember the words as best they could for a later memory test. On the single-item trials, participants were instructed to hit either “z” to respond “not studied” or “/” to respond “studied” (on a standard QWERTY keyboard), and they were asked to keep their index fingers on the “z” and “/” key throughout the test. On the forced-choice trials, participants were asked to respond “z” or “/” to indicate that the word on the left side or the right side was the studied word, respectively.

Throughout the experiment, we used points to reinforce desired performance patterns, and participants were asked to earn as many points as possible (points were not linked to monetary payments). To introduce biases in single-item decisions, we designated one “safe” response that paid 3 points for correct responses and deducted only 1 point for errors and one “risky” response that paid 1 point for corrects and deducted 3 points for errors. To demonstrate that it was better to use the “safe” than the “risky” response when errors were likely, the first session began with an activity in which participants had to choose between “Studied” and “Not Studied” responses with no way to determine which one was correct. They were told that they just had to guess, and that they could only expect to be correct half of the time. In one round of 8 trials, they were asked to choose the risky response, which was indicated by red font for the response label, on every trial. The number of points gained or lost appeared on the screen for .5 s after each response, and after all responses they saw feedback on the total points gained or lost. If they followed instructions and used the risky response every time, then they lost 8 points total. That is, each of the 4 incorrect responses lost 3 points and each of the 4 correct responses earned 1 point, $4 \times -3 + 4 \times 1 = -8$. On a second round of 8 trials, they were asked to choose the safe response, indicated with green font, on every trial. If they followed instructions, they earned a total of 8 points, $4 \times 3 + 4 \times -1 = 8$. After both cycles, the experimenter summarized the purpose of

the activity by noting that participants should use the safe response when they are uncertain that they can respond correctly to avoid losing a lot of points.

The test had blocks of liberal and conservative payoff schemes. Each block comprised 12 single-item test trials, and the first block also included the four filler trials at the beginning of test list. Blocks were randomly assigned to be liberal or conservative with the constraint that there were an equal number of blocks in each condition. In liberal blocks, “studied” responses were safe and “not studied” responses were risky (by the definition listed above). In conservative blocks, “not studied” responses were safe and “studied” responses were risky. Throughout the test, safe responses appeared in green font and risky responses appeared in red font. For example, on conservative blocks “Not Studied” appeared in green and “Studied” appeared in red to label the alternative responses on each trial. Before every block, participants were notified if the “Studied” or “Not Studied” response would be the safe response in the following block, and the initial test instructions reviewed what “safe” and “risky” meant. Immediate point feedback was *not* provided for single-item trials, because this would allow participants to select correct responses on later forced-choice trials by remembering the feedback. However, participants still earned or lost points depending on whether they chose the safe or risky response and whether their response was correct, and they saw feedback on point totals after each block (described below).

To ensure that participants attended to the bias information during the single-item test, they were asked to make two responses on each trial: a guess before seeing the test word and then a “real” response after seeing it. At the beginning of each single-item trial, the “studied” and “not studied” choices appeared on the screen before the test word was presented, with the safe response for that block in green and the risky response in red. Participants had to guess which

would be the correct response for that trial, and they were advised that they should choose the safe response in their initial guess to avoid losing a lot of points for wrong guesses. These initial responses not only focused participants on the bias information, but also served as a manipulation check to eliminate participants who did not understand the “safe” and “risky” distinction or did not respond to it in a reasonable manner. After the first response, the test word appeared and the participant had a second opportunity to respond based on both their memory for the word and their knowledge of the payoff scheme. When participants lost points overall for their initial guesses in a block, they saw an extra feedback screen after the block reporting the number of points they lost for their guesses and reminding them that the better strategy was to guess the safe response when they didn’t know the correct answer.

Forced-choice trials were interspersed with the single-item trials for each test list. Specifically, after each block of 12 single-item trials,⁴ the computer created forced-choice trials by pairing targets and lures that received the same single-item response, meaning that every trial had one word that was previously classified correctly and one that was previously classified incorrectly (see Table 1). For example, a lure that was falsely recognized as “studied” would be randomly paired with a target that was correctly called “studied.” Participants were informed that every forced-choice trial would have either two words they called “studied” or two words they called “not studied,” but that in reality one of the words was studied and one was not. They were asked to indicate which of the two words was actually studied. Words in the forced-choice trials always came from the most recent single-item block, so the lag between an item’s appearance in the single-item and forced-choice trials was on the order of seconds. The computer created as many forced-choice trials as possible after each block based on the participant’s responses, so the

⁴ The filler trials did not count for the block divisions, so the first block appeared to have 16 words from the participant’s perspective, but the four fillers were not eligible to appear as forced-choice trials.

number of forced-choice trials varied from one participant to the next. Some blocks could have no forced-choice trials, for example if the participant did not make any errors, and in this case participants simply hit the spacebar to advance after finishing the single-item trials for the block.

On the forced-choice trials, participants earned 2 points for correct responses and lost 2 points for incorrect responses, and a screen showing either +2 or -2 appeared after each response to provide immediate feedback. The point feedback remained on the screen for one second. After all forced-choice trials were completed, the block ended with a feedback screen reporting the number of points gained or lost for the initial guesses and second responses on single-word trials and for the forced-choice trials (identified to the participants as “two-word” trials) as well as the total points earned so far in the experiment. For example, the screen might read:

This block:

Single-word initial guesses: 36 points

Single-word 2nd response: 48 points

Two-word trials: 12 points

Experiment total: 128 points

This screen remained visible until the participant pressed SPACE.

Results

The critical analyses reported herein were described in the preregistration documents on the project’s OSF page (review link:

https://osf.io/kn2tw/?view_only=7c207f36faf745fe85ff47acc75ed648). The analyses address

three goals: (1) evaluate single-item performance to ensure that the bias manipulation was effective and the data conform to typical response patterns, (2) evaluate the effect of the bias manipulation on forced-choice accuracy, and (3) define and test forced-choice predictions of the

2HT and UVSD models for the specific parameter values estimated from the single-item trials. We used Bayesian parameter estimation for inferential analyses, and we report the median (*Md*) and a 90% credibility interval (*CI*) to summarize the marginal posterior distribution for each parameter of interest. The interval end points are the .05 and .95 quantile of the marginal posterior distribution. For interested readers, the Appendix reports the results of null hypothesis significance tests for key comparisons.

Single-Item Recognition

Table 2 shows the average hit and false-alarm rates from the single-item trials. The bias manipulation appeared to have the intended effect, as participants were more likely to respond “studied” for both targets and lures in the liberal condition than in the conservative condition. Moreover, the bias shifts were relatively large: we observed a .16 difference in false-alarm rate and a .13 difference in hit rate between the bias conditions. These large effects are advantageous for testing the model predictions (see Figure 3), but reviewers were concerned that our “enforced guessing” procedure – i.e., having participants respond based on the payoff information before seeing the test word and then respond again after the word appeared – encouraged idiosyncratic guessing strategies. Notably, some participants may have been likely to repeat their initial guess after the test word appeared without attempting to remember the word. We conducted a series of analyses to explore this possibility.⁵

First, we compared overall single-item recognition performance to a recent study that was very similar to the current study but did not involve a bias manipulation or the enforced guessing procedure (Starns et al., 2018, Experiment 2). Performance levels were similar. The comparison study had a hit rate (HR) of .64 and a false-alarm rate (FAR) of .26 compared to HR = .63 and

⁵ These analyses were conducted on the 160 participants who were eligible for analyses based on the pre-registered exclusion criteria.

FAR = .17 in the current conservative condition and HR = .76 and FAR = .33 in the current liberal condition. Thus, performance levels with the enforced guessing procedure are in the expected range, suggesting that participants made a normal attempt to remember the test items.

Second, we evaluated the response time (RT) data to check for fast guessing. Figure 6 shows the overall proportion correct across all single-item recognition responses for a series of 100 ms RT bins with the bin centers ranging from 100-3000 ms. Performance is at chance levels in the earliest bins, rises to just above chance for the 350-450 and the 450-550 ms bins, then dramatically increases for the 550-650 ms bin. Based on this plot, we set 550 as the cutoff for fast guesses. This is later than a typical fast-guess cutoff for recognition (e.g., Ratcliff, Thapar, & McKoon, 2003), but our goal was to exclude as many guess responses as possible even if we also exclude a few fast, memory-based responses. By this definition, only 3% of responses were classified as fast guesses. Notably, 90% of the RTs for the initial guess responses fell below our criterion for fast guesses, so we can be confident that the vast majority of fast guesses for the subsequent recognition decisions would be identified as such. In summary, the RT data are consistent with a normal recognition process and show no sign that our enforced guessing procedure encouraged participants to repeat their guess response without trying to remember the test item.

To ensure that guessing strategies did not influence the results, our analyses exclude any participants who made over 10% of their single-item recognition responses in less than 550 ms. This eliminated 13 participants, and with these exclusions the proportion of trials with RTs below 550 ms fell to 0.6%. These additional exclusions had very little impact on single-item performance measures; for example, the cells in Table 1 go from [.76, .33, .63, .17] before the guessing exclusions to [.76, .31, .64, .17] after.

Single-Item Modeling Methods

This section describes our Bayesian models for single-item performance in detail. Readers who are not interested in these details can safely skip it. We analyzed the single-item recognition data by applying hierarchical Bayesian versions of the 2HT and UVSD models using JAGS software (Plummer, 2003). Modeling and analysis code files are available on the project's OSF page (review link: https://osf.io/kn2tw/?view_only=7c207f36faf745fe85ff47acc75ed648). The data from each participant has 4 degrees of freedom (*df*), one each from targets and lures in the liberal and conservative bias conditions (performance in each of these conditions is defined by two response frequencies that must sum to the total number of trials, thus representing one *df*).

For the 2HT model, a given participant's data were predicted based on 4 parameters: *Do* (the probability of detecting that a target was studied or "old"), *Dn* (the probability of detecting that a lure was not studied or "new"), *gc* (the probability of guessing "old" in the conservative condition), and *gl* (the probability of guessing "old" in the liberal condition). The probability of responding "old" for targets and lures in the conservative and liberal conditions was calculated from these parameters with the standard model equations (see the Introduction for these equations). In the JAGS model, each test trial was modeled as a Bernoulli distribution with the probability parameter defined by the model predictions. The participant-level model parameters were converted to a probit scale, and $\text{probit}(gc)$ and $\text{probit}(gl)$ were re-parameterized as the average probit across both bias conditions, *gav*, and the distance between $\text{probit}(gc)$ and $\text{probit}(gl)$, or the bias effect, *be*. Every participant was allowed a unique value of these four probit-scale parameters, and we used Gaussian distributions to represent variation in each parameter from one participant to the next. Each across-participant distribution was defined by a

mean and standard deviation parameter, so there were 8 hyperparameters in total for the four participant-level parameters. We used loosely informative priors on the hyperparameters defining the across-participant distributions. For the mean of the across-participant distributions of $\text{probit}(Do)$ and $\text{probit}(Dn)$, the prior distribution was a Gaussian with a mean of zero and a standard deviation of 0.75. For the mean of the across-participant distributions of gav and be , the prior distribution was a Gaussian with a mean of zero and a standard deviation of 0.5. For the hyperparameters defining the standard deviation of each of the across-participant distributions, the prior distribution was a uniform distribution from 0.05 to 2.

Our analyses focused on the parameters defining the across-participant distributions. Specifically, we characterized the memory performance of a typical participant by evaluating the posterior distributions for the mean of the across-participant distribution of $\text{probit}(Do)$ and $\text{probit}(Dn)$, and we converted all samples from these posterior distributions back to the probability scale for ease of interpretation. To characterize the decision biases of a typical participant, we created posterior distributions for the across-participant mean of $\text{probit}(gc)$ and $\text{probit}(gl)$ by transforming the posterior distributions for the means of the across-participant distributions of the gav and be parameters [$\text{probit}(gc) = gav - .5 \times be$ and $\text{probit}(gl) = gav + .5 \times be$]. Again, we converted all samples from these posterior distributions to the probability scale for ease of interpretation.

For the UVSD model, a given participant's data were predicted based on 4 parameters: μ (the average memory strength for targets), σ (the standard deviation in memory strength for targets), c (the average criterion across the two bias conditions measured as a distance from $.5 \times \mu$), and be (the bias effect, defined as the distance from cc to cl , which are the criteria in the conservative and liberal bias conditions, respectively). As an example, if $c = .2$ and $\mu = 2$, then

the absolute position of the average criterion is 1.2. Further, if $be = -.4$, then $cc = 1.4$ and $cl = 1$.

The probability of responding “old” for targets and lures in the conservative and liberal conditions was calculated from these parameters with the standard model equations (Macmillan & Creelman, 2005). In the JAGS model, each test trial was modeled as a Bernoulli distribution with the probability parameter defined by the model predictions. Every participant had a unique set of model parameters, and we used Gaussian distributions to represent across-participant variation in μ , c , be , and $\log(\sigma)$. Each across-participant distribution was defined by a mean and standard deviation parameter, so there were 8 hyperparameters in total for the four participant-level parameters. We used loosely informative prior distributions on the hyperparameters. For the average μ value across participants, the prior distribution was a truncated Gaussian with a mean of 1, a standard deviation of 2, and a lower limit of 0. For the average c and be values across participants, the prior distributions were Gaussians with a mean of 0 and a standard deviation of 1. For the average $\log(\sigma)$ across participants, the prior was a Gaussian distribution with a mean of 0 and a standard deviation of 0.5. The priors on the standard deviation parameters for the across-participant distributions were uniform distributions from 0.05 to 2.

As with the 2HT model, our analyses for the UVSD model focused on the parameters defining the across-participant distributions. Specifically, we characterized the memory performance of a typical participant by evaluating the posterior distributions for the across-participant mean of the μ and $\log(\sigma)$ parameters. To aid interpretation, we converted the latter back to the raw scale. To characterize the decision biases of a typical participant, we created posterior distributions for the across-participant mean of the cc and cl parameters by transforming the posterior distributions for the mean c and mean be parameters [$cc = c + .5 \times be$ and $cl = c - .5 \times be$].

Single-Item Recognition Modeling Results

For the 2HT model, results showed that the probability of detecting a target as old (posterior $Md = .47$, 90% $CI [.43, .51]$) was higher than the probability of detecting a lure as new (posterior $Md = .34$, 90% $CI [.29, .41]$). The overall probability of responding “studied” when guessing was slightly below .5 (posterior $Md = .40$, 90% $CI [.37, .43]$), indicating that participants tended to adopt a conservative response strategy overall. However, the payoff manipulation had a clear effect on guessing strategies, such that the probability of guessing “studied” was higher in the liberal condition (posterior $Md = .53$, 90% $CI [.49, .56]$) than the conservative condition (posterior $Md = .28$, 90% $CI [.25, .31]$).

For the UVSD model, results showed that a typical participant had a target strength distribution with a mean value of 1.35 (90% $CI [1.12, 1.52]$) and a standard deviation of 1.30 (90% $CI [1.22, 1.39]$). The latter value corresponds to a ratio of lure-to-target standard deviation values of 0.77 (90% $CI [0.72, 0.82]$), so these data are consistent with the “benchmark” value of 0.8 for this ratio (Ratcliff et al., 1992; Wixted, 2007). Results also showed a large shift in the response criterion from the liberal condition (posterior $Md = -0.25$, 90% $CI [-0.31, -0.18]$) to the conservative condition (posterior $Md = 0.28$, 90% $CI [0.22, 0.34]$).

In summary, both models characterized the data in ways that are quite similar to what is typically found in the literature at large, with reasonable levels of memory performance and clear response-bias effects in response to the payoff manipulation (e.g., Bröder & Schütz, 2009; Dube & Rotello, 2012).

Forced-choice Accuracy

Table 2 shows accuracy on the forced-choice trials. Average accuracy levels were much higher on S-S trials than on N-N trials, a finding that replicates earlier investigations using this “error correction” testing paradigm (Starns et al., 2018; Voormann et al., 2021). This finding indicates that participants are very capable of distinguishing a target that they correctly recognized from a lure that they falsely recognized, but they struggle to distinguish a target that they falsely rejected from a lure that they correctly rejected. However, accuracy was very stable across the two bias conditions, which is notable when one considers the large bias effects on single-item performance (see Table 1).

Forced-choice Modeling Methods

This section reports the details of our Bayesian analyses for forced choice performance. Readers who are not interested in these details can safely skip it. We submitted the forced-choice data to a hierarchical Bayesian logistic regression using trial type (N-N versus S-S) and payoff bias (liberal versus conservative) as dichotomous predictors. Modeling and analysis code files are available on the project’s OSF page (review link: https://osf.io/kn2tw/?view_only=7c207f36faf745fe85ff47acc75ed648). In the JAGS model, the accuracy of each forced-choice trial was modeled as a Bernoulli distribution with the probability parameter determined by the logistic equation. The participant-level logistic parameters included a grand mean, gm , representing the logit-transformed overall probability correct; a trial-type effect, te , representing the difference between the logit-transformed probability correct for N-N and S-S trials; a bias effect, be , representing the difference between the logit-transformed probability correct for the liberal and conservative conditions; and an interaction effect, ie , representing the difference between the bias effect on N-N and S-S trials. Each participant had unique values of the four logistic parameters, and across-participant variation in parameters was

modeled with Gaussian distributions. Each across-participant distribution was characterized by mean and standard deviation parameters, so there were eight hyperparameters. We used loosely informative priors on these hyperparameters; specifically, for each of the logistic parameters, the prior distribution on the across-participant mean was a Gaussian distribution with a mean of 0 and a standard deviation of 1 and the prior distribution on the across-participant standard deviation was a uniform distribution from 0.05 to 2.

Our analyses focused on average performance levels across participants. We defined posterior distributions for the logit-transformed probability correct in each of the four conditions and then converted all posterior samples back to the probability scale to facilitate interpretation. All reported analyses were performed on these probability-scale samples.

Forced-choice Modeling Results

The logistic model showed a clear indication that forced-choice percent correct was higher for the S-S trials ($Md = .77$, 90% CI [.76, .78]) than N-N trials ($Md = .63$, 90% CI [.61, .64]), and the credible interval on the posterior distribution of differences indicated an effect of over 10 percentage points (90% CI [.13, .16]). The model found slightly higher performance for the conservative condition ($Md = .71$, 90% CI [.70, .73]) than the liberal condition ($Md = .69$, 90% CI [.67, .70]), but the credible interval on the difference reveals that this is a very small decrement, at best (90% CI [.01, .04]). The posterior distribution for the interaction between trial type and bias showed that there was most likely a very small interaction effect ($Md = .03$, 90% CI [-.002, .05]), although results did not strongly rule out null or functionally-null interactions. Recall that the interaction term, or Difference of Differences (DoD), is the difference between bias effects in S-S and N-N trials. Both trial types had higher accuracy in conservative than liberal blocks, with a small difference for S-S trials ($Md = .04$, 90% CI [.02, .06]) and a very

small difference for N-N trials ($Md = .01$, 90% $CI [-.01, .03]$). To summarize, the model results are consistent with a small bias effect that is slightly larger for S-S than N-N trials, but the population effects could be very small or even functionally null.

Model Testing

The previous section revealed that the forced-choice results did not show a clear signal of bias effects on forced-choice accuracy, but the size of the effect predicted by the 2HT and UVSD models varies considerably across parameter sets. The preregistration document for this project (review link: https://osf.io/kn2tw/?view_only=7c207f36faf745fe85ff47acc75ed648) outlines a strategy for testing each model's specific quantitative predictions for the current data set, and this section reports the results of those tests. Essentially, the strategy involved defining likely parameter values by modeling data from the single-item trials and generating predicted forced-choice performance with those tailored parameters. The models did not “see” any of the forced-choice data, so an ability to match these data would indicate true predictive value. This model selection strategy has been called the “validity generalization criterion” method (Mosier, 1951), and Busemeyer and Wang (2000) provided a clear summary of its virtues for testing alternative models of cognition.

Predictions were defined by running 5000 replications of the full experiment with model-simulated data. The simulated experiments matched all methodological details of the real experiments, including the total number of participants and the number of forced-choice trials completed by each participant. The method for generating model-simulated data preserved all levels of uncertainty about forced-choice performance, including sampling variability and uncertainty in the population value of the model parameters for each participant, with the latter represented by the participant-level posterior distributions from the hierarchical Bayesian model

fits to the single-item data. We computed three target values for each replication: (1) the average bias effect for S-S trials, defined as the average percent correct for S-S trials in the conservative payoff condition minus the liberal payoff condition; (2) the average bias effect for N-N trials; and (3) the average interaction effect or Difference of Differences (DoD), defined as the average bias effect for S-S trials minus the average bias effect for N-N trials. The distributions of these values across the 5000 replications defined the prediction distributions.

In detail, for each decision model (2HT and UVSD), we used the following steps for each simulated experiment: (1) we randomly selected one sample from the joint posterior distribution that was produced by the hierarchical Bayesian fits; (2) we took the participant-level parameters from the selected posterior sample, which yielded a unique parameter set for each participant; (3) we evaluated the predicted percent correct for each simulated participant in the four conditions (liberal and conservative S-S and N-N trials) using the sampled parameter values and the model's prediction equations (see the Introduction for these equations); (4) we randomly sampled forced-choice trial outcomes (correct or error) from a binomial distribution for each condition within each participant based on the predicted probability correct and the number of forced-choice trials for that participant in that condition, with each full set of trial outcomes across conditions representing a *simulated* participant; (5) we divided the sampled number of correct trials by the total number of trials for each condition within each simulated participant to get accuracy rates; (6) we averaged accuracy values across simulated participants in each of the four conditions⁶; and (7) we used these average accuracy values to calculate the three effects of

⁶ Although trial counts could vary across participants, we took simple averages, not weighted averages. In the prediction simulations for the preregistration, we also tried simulations that took weighted averages of percent correct in each of the four conditions (S-S and N-N trials on liberal and conservative blocks) and computed the interaction effect with these weighted averages, and the results were very similar. We decided to use unweighted averages for the simulation in the pre-registration (and the one reported herein) since this will give the same interaction effect as a simple ANOVA on the percent correct data.

interest, namely the bias effect on S-S trials, the bias effect on N-N trials, and the interaction of bias and trial type. This procedure was repeated 5000 times for each decision model to define prediction distributions for the effects. To determine responses on the forced choice trials, we retained the memory state of each item from the single-item test. For the 2HT model, this means that items that could (not) be detected on the single-item trials also could (not) be detected on the forced-choice trials. For the UV model, this means that items had the same strength on the single-item and forced-choice tests. This is a simplifying assumption, but with blocks of only 12 items memory states should be very consistent between the two trial types. Later on, we will consider whether changing memory states could alter model predictions.

Figure 7 shows the prediction distributions for the S-S bias effect, N-N bias effect, and the DoD, with vertical lines marking the observed results. The prediction distributions show the same pattern noted in the Introduction. The 2HT model predicts a positive bias effect on S-S trials (i.e., greater accuracy for the conservative than liberal condition), a negative effect for N-N trials, and a positive DoD. In qualitative terms, the 2HT correctly predicts the positive effect observed in S-S trials as well as the positive DoD, but fails to capture the observed positive effect in N-N trials. In turn, the UVSD makes qualitative misses at the level of S-S accuracy and DoD, but correctly predicts the positive bias effect at the level of N-N trials. Quantitatively speaking, both models fail, in the sense that their predictions are considerably more extreme than the observed effects.

Overall, we see that each model fails in some way or another. This state of affairs introduces the challenge of finding a model that could provide an alternative account that is satisfactory at a qualitative as well as a quantitative level. In the sections below, we analyze a number of alternative models that replace or relax some of the assumptions made by the 2HT and

UVSD models. The goal here is to try to find among the candidates at least one that can capture the qualitative pattern and then subject it a predictive test by challenging the model to match the forced-choice data using parameter values informed by the single-item data.

Exploring Model Variants

As previously stated, one of the assumptions made when obtaining predictions from the models is that the mental states underlying single-item recognition judgments carry on to the forced-choice judgments. It is certainly plausible that memory for a given item would be highly consistent across the two trial types given that we relied on short single-item recognition blocks, which meant that the forced-choice judgments were conducted shortly after the single-item judgments. Nevertheless, it is important to consider whether changes in memory would change model predictions for our paradigm, and we explore this question with a series of model simulations.

Before considering quantitative results, it is important to note that the predictions we outline in the introduction are based on “deep” model mechanisms. The predictions for both models are based on selection effects introduced by biased responding in the single-item trials. A core assumption of the 2HT model is that response biases are based on different probabilistic response patterns from an uncertainty state. Thus, single-item bias affects the proportion of correct “old” and “new” responses that are based on memory versus lucky guesses. Compared to neutral responding, liberal bias mixes in more lucky guesses in the pool of correct “old” responses and conservative bias mixes in more lucky guesses in the pool of correct “new” responses. Mixing in more failed-detection items impairs forced-choice accuracy. A core assumption of the UVSD model is that response biases are produced by changing the level of memory strength needed to justify an “old” response. Thus, single-item bias affects the

distribution of memory strength values conditioned on certain responses. For example, with conservative responding, all items called “old” will have high memory strength, even lures called “old.” Extreme bias “bunches up” the conditional evidence distributions and makes it difficult to distinguish targets and lures that received the same single-item response (Figures 4 and 5). These selection effects are based on core model mechanisms, and eliminating them completely would require complete independence between retrieval outcomes on the single-item and forced-choice trials. For example, participants would need to commonly remember items that they failed to remember a few seconds ago; in fact, they would need to be equally likely to remember these items as items that they successfully remembered a few seconds ago. Thus, intuition suggests that model predictions should be robust when various task-specific processes are added to the models. The simulations below explore whether this intuition is sound.

Allowing for Changes in Memory States

Here, we consider how the 2HT model predictions would be affected if we allowed the opportunity for non-detected items from the single-item test to be detected when presented in a forced-choice trial (for a relevant discussion, see Malejka & Bröder, 2016).

In order to test this possibility, we developed an extended 2HT model, which we will refer to as the 2HT* model, that allows for non-detected items to be subsequently detected with probabilities D_o^* and D_n^* .⁷ The predictions made by this model are illustrated in Figure 8. What these predictions show is that the increase of D_o^* and D_n^* does not affect the general prediction pattern: The model still predicts a positive bias effect in S-S trials, a negative effect in N-N trials, and a positive DoD. Based on this example, it is clear that relaxing the assumption that memory

⁷ Model equations can be found in the Supplemental Materials.

states are preserved across single-item and forced-choice recognition trials does not salvage the 2HT model.

Within-Item Memory Strength Variability

We also considered an analogous modification to the UVSD model, which we will refer to as the UVSD* model. Specifically, we assumed that the memory-strength value associated with each test item can vary between the single-item and the forced-choice trials in which they are presented. Specifically, we assumed that each memory-strength value encountered in a single-item trial was subsequently distorted by adding to it an independent noise term ε , which followed a Gaussian distribution with mean 0 and standard deviation σ_ε . The predictions made by the UVSD* model, which are illustrated in Figure 8, show that the introduction of noise does not change predictions in a qualitative way – it only decreases overall performance.⁸ As a final check, we also considered a variant of the UVSD* model in which the noise term ε is strictly positive, capturing the notion that memory-strength is boosted by the initial single-item test phase. This variant also made predictions that are qualitatively similar to the UVSD model.

Criterion Noise

In the context of SDT modeling, it is typically assumed that the response criterion κ is fixed across trials. The adequacy of this assumption is a somewhat contentious point in the context of recognition memory, with some researchers suggesting that it is more appropriate to assume that κ varies across test trials (see Benjamin, 2013; Benjamin, Diaz, & Wee, 2008; Benjamin, Tullis, & Lee, 2013; Kellen, Klauer, & Singmann, 2012; 2013). In any case, it is interesting to check whether the presence of criterion noise could change the UVSD predictions

⁸ Because these predictions were obtained via simulation, the predicted values tend to be somewhat noisy, especially in the case of extreme response criteria. We therefore limited criteria to the [-2,2] interval. For further details on this simulation, see the Supplemental Materials.

in the direction of the observed effects. To evaluate this conjecture, we established the UVSD_{CN} model, which assumes that κ follows a Gaussian distribution with mean μ_κ and standard deviation σ_κ . In a simulation, each memory strength was compared with an independently-sampled criterion. As shown in Figure 8, the presence of considerable levels of criterion noise ($\sigma_\kappa = 0.75$) can affect the direction of the predictions in S-S trials, such that accuracy is no longer strictly decreasing as the (average) criterion becomes more conservative. However, this effect does not change the prediction of negative DoDs. Although somewhat promising, these results should be taken with a grain of salt. After all, criterion noise is believed to be most prevalent among *confidence criteria*, not the criterion that determines the binary “old”/“new” judgment (e.g., Mueller & Weidemann, 2012; see also Benjamin et al., 2013). One of the reasons for such a belief is that the presence of relatively large magnitudes of criterion noise would lead to gross mispredictions across recognition-memory tasks, which turn out to not occur (e.g., Jang, Wixted, & Huber, 2009; Kellen et al., 2012; 2021).

Exploring Additional Decision Models

Our original goal was to contrast the 2HT and UVSD models, but our results are also relevant for alternative approaches. In this section, we explore two other popular models of recognition tasks.

Low-Threshold Model

Although it was not our original intention, it turns out that the “error correction” paradigm used in the present work establishes a strong test for another candidate model in the literature, the *Low-Threshold Model* (Luce, 1963; see also Kellen, Erdfelder, Malmberg, Dube, & Criss, 2016). Similar to the 2HT model, the low-threshold model assumes that responses consist of a mixture of detection and guessing-based judgments. But in contrast with the 2HT

model, it assumes that both targets and lures can reach the Detect-Old state (with probabilities D_t and D_l), and precludes the existence of a Detect-New state. Finally, it assumes that responses are produced under “liberal” and “conservative” response policies. According the former, the probability of an “old” response for targets and lures correspond to:

$$P(\text{"old"}|\text{target}) = D_t + (1 - D_t)g$$

$$P(\text{"old"}|\text{lure}) = D_l + (1 - D_l)g$$

whereas under a “conservative” policy, they are given by:

$$P(\text{"old"}|\text{target}) = D_t g'$$

$$P(\text{"old"}|\text{lure}) = D_l g'$$

where g and g' are guessing parameters. Specifically, under the liberal policy, participants respond “old” for every item in the Detect-Old state and also for a proportion g of items that are not in the Detect-Old state. Under the conservative policy, participants respond “new” for all items that are not in the Detect-Old state and also for a proportion g' of items that are in the Detect-Old state.

The low-threshold model makes a clear prediction, namely that forced-choice accuracy for any given participant should be at chance for at least one of the forced-choice trial types (S-S and N-N). If a participant adopts a “liberal” policy, then all “new” judgments are made to items that did not produce the Detect-Old state, and this is true for both targets and lures called “new.” This implies that accuracy in N-N trials should at chance level (i.e., .50). Alternatively, if a participant adopts a “conservative” policy, then performance in S-S trials should be at chance by the same logic.

We tested these predictions by checking participants’ performance in both N-N and S-S trials. Figure 9 shows the estimates (along with 90% credibility intervals) of individual accuracy

obtained from our Bayesian analysis of the forced-choice data. We classified performance as above chance if the posterior distribution indicated at least a 95% probability that percent correct was over .5 for that participant and trial type (S-S or N-N). By this standard, 106 of 147 participants performed above chance on *both* trial types in the conservative bias condition and 98 did so in the liberal bias condition. These results provide clear evidence that it is possible to accurately discriminate targets from lures on both trial types, a problematic result for the low-threshold model.

The low-threshold model could potentially be altered to accommodate this finding under the assumption that individual participants alternate between the liberal and conservative response strategies across trials. We will leave this for future studies to explore, but we are skeptical that the model could produce forced-choice performance levels that match empirical results even under these less restrictive assumptions.

Dual-Process Signal Detection Model

Finally, we considered the *Dual Process Signal Detection* (DPSD) model (Mandler, 1980; Yonelinas, 1997), which combines a threshold process termed “recollection” along with an evaluation of memory strengths that is equivalent to a Gaussian signal detection model (when assuming that $\sigma_t = \sigma_l = 1$). Both processes are assumed to be independent. When judging a target item, an individual can recollect it with probability R . When recollection fails, with probability $1 - R$, the recognition judgment is based on the memory strength associated with that target. In turn, lures are always evaluated in terms of their memory strength.

According to the DPSD model, the probability of a correct response in S-S and N-N trials corresponds to:

$$P(\text{correct}, S - S) = \frac{R(1 - F(\kappa)) + (1 - R) \int_{\kappa}^{\infty} f(x, \mu_t, 1)(F(x) - F(\kappa))dx}{R(1 - F(\kappa)) + (1 - R)(1 - F(\kappa, \mu_t, 1))(1 - F(\kappa))}$$

$$P(\text{correct}, N - N) = \int_{-\infty}^{\kappa} \frac{f(x, \mu_t, 1) F(x)}{F(\kappa, \mu_t, 1) F(\kappa)} dx$$

Figure 10 illustrates the DPSD's predictions for the S-S and N-N trials. For the S-S trials, predictions with $R=0$ show the same qualitative pattern as the UVSD model, with decreasing forced-choice accuracy as the criterion value increases. However, if recollection occurs with some non-zero probability, the model often predicts that accuracy increases as the response criterion becomes more conservative. The reason behind this shift is that recollected targets will always be correctly selected in a S-S trial, *regardless of the response criterion*. As the response criterion becomes increasingly conservative, the relative proportion of recollected targets in S-S trials will increase, which in turn will increase accuracy. The effect of recollection is also reflected in the DoDs, which can be either positive or negative.

Based on this first evaluation, the DPSD appears to be suitable *candidate* for describing the forced-choice data collected in our experiments. After all, this apparent suitability might be merely due to the fact that this appears to be a somewhat flexible model. In order to evaluate the DPSD, we used the same predictive testing approach adopted earlier with the 2HT and UVSD models. Specifically, we fitted the DPSD to the single-item judgments and used the parameter estimates to predict accuracy in the forced-choice judgments. That is, we fit this model to the single-item recognition data and then ran model simulations to define what it predicts for the forced-choice data.

We fit the DPSD model to the single-item recognition data using a hierarchical Bayesian approach that was very similar to the 2HT and UVSD fits above. Each participant was allowed unique values of four DPSD parameters: the probability of target recollection (R), the distance between the target and lure familiarity distributions (d'), the overall response criterion measured

as a distance from the midpoint between the lure and target distributions (c), and the bias effect, operationalized as distance between the criteria for the liberal and conservative bias conditions (be). Thus, a given participant's response criterion was $d'/2 + c - .5 \times be$ for the liberal condition and $d'/2 + c - .5 \times be$ for the conservative condition. We represented across-participant variation in parameter values with Gaussian distributions, so there were eight hyperparameters for the mean and standard deviation of the across-participant distributions of d' , c , be , and the probit of R . We used loosely informative priors on the hyperparameters. For the means of the across-participant distributions, we used a truncated normal with a mean of 1, a standard deviation of 2, and a lower bound of 0 for d' ; a Gaussian with a mean of 0 and a standard deviation of 1 for c ; a Gaussian with a mean of zero and a standard deviation of 1 for be ; and a Gaussian with a mean of 0 and a standard deviation of 0.75 for $\text{probit}(R)$. For the standard deviations of the across-participant distributions, we used a uniform distribution from 0.05 to 2 for d' , c , and be ; and a uniform distribution from 0.05 to 0.75 for $\text{probit}(R)$.

Posterior distributions on the means of the across-participant distributions indicated a moderate probability of recollection (R ; $Md = .28$, 90% CI [.22, .33]) and relatively low levels of familiarity-based discriminability (d' ; $Md = 0.27$, 90% CI [0.03, 0.68]). The familiarity response criterion was lower in the liberal condition ($Md = 0.06$, 90% CI [-0.01, 0.13]) than the conservative condition ($Md = 0.53$, 90% CI [0.45, 0.61]).⁹

We defined predictions for the forced-choice data using the same simulation method described for the UV and 2HT models. For S-S trials, we assumed that participants responded correctly if the target was recollected or if the target was not recollected but had a higher familiarity value than the lure, and that participants responded incorrectly when the target was

⁹ These values are in reference to the midpoint between the lure and target familiarity distributions.

not recollected and had a lower familiarity value than the lure. For N-N trials, we assumed that participants responded correctly if the target had a higher familiarity value than the lure and responded incorrectly otherwise. Figure 11 shows distributions of the DPSD predictions for average effects across participants, with vertical lines marking the observed values. The model predicts small negative bias effects (i.e., higher forced-choice accuracy in the conservative than the liberal condition) for both S-S and N-N trials and a small positive interaction. The data are generally in good accord with these predictions. The observed S-S effect is fairly far out on the left tail of the prediction distribution, but the miss is much less severe than the misses for the UVSD and 2HT models (see Figure 7). Note that allowing for small changes in memory from the single-item to forced-choice trials will pull the prediction distributions a bit closer to zero, resulting in even better predictions. When all effects are considered, the DPSD model predicted forced-choice performance much better than the other models that we considered.

General Discussion

This study investigated the mechanism of recognition memory with experiments designed to test two models that are often used to interpret recognition results, the two high threshold (2HT) model and the unequal variance signal detection (UVSD) model. Previous studies have tested these models extensively with techniques such as ROC functions, ranking tasks, and RT modeling (for recent reviews, see Pazzaglia, Dube, & Rotello, 2013, and the response by Batchelder & Alexander, 2013). Our study tested these models with single-item recognition followed by forced-choice recognition on the same items. Although a number previous studies have used the forced-choice testing method (Province, & Rouder, 2012; Kellen, Singmann, Vogt, & Klauer, 2015; Kroll, Yonelinas, Dobbins, & Frederick, 2002; Parks & Yonelinas, 2009; Kellen & Klauer, 2011), or have even conducted single-item and forced-choice testing

simultaneously to discriminate the two recognition memory models (Jang, Wixted, & Huber, 2009; Smith & Duncan, 2004), our experiment provided a new test of the models by evaluating the effect of response biases on the probability that participants could correct their errors from earlier single-item recognition trials when the words later appeared in forced-choice trials.

When combined with a bias manipulation, this “error correction” paradigm places very strong constraints on the models, as illustrated in Figure 3. Both models predict that inducing response biases for the single-item trials should have consequences for accuracy on subsequent forced-choice trials. The 2HT model must predict that more conservative single-item responding increases forced-choice accuracy on trials that have a target and a lure that were both previously called “studied” (S-S trials) and decreases forced-choice accuracy on trials with a target and a lure that were both called “not studied” (N-N trials). Thus, it also must predict a positive interaction, or a Difference of Differences (DoD) when the S-S effect is subtracted from the N-N effect. The UVSD model predicts that more conservative single-item responding decreases forced-choice accuracy on S-S trials and increases accuracy on N-N trials. When these effects are combined, the UVSD model predicts a negative interaction under any plausible parameterizations.

Fitting the 2HT and UVSD models to the single-item data revealed that our payoff manipulation induced strong response biases, so the models predict large effects on forced-choice accuracy. Nevertheless, results showed that the bias effect on forced-choice accuracy was near zero for both N-N and S-S trials, as was the interaction of trial type and bias. Prediction simulations informed by the single-item parameter estimates verified that the models predicted clear bias effects on forced-choice responding, and the observed near-zero effects were well out of range of the prediction distributions for either model.

Our study adds to previous efforts to test decision models with joint single-item and forced-choice recognition data (Jang et al., 2009; Kellen, Winiger, Dunn, & Singmann, 2019; Smith, & Duncan, 2004). These studies explored a variety of different models and theoretical issues, but one common theme is that all of these studies included the UVSD model, and it consistently provided a good fit to the joint single-item and forced-choice data (or made good predictions about one form of data based on the other). This might seem to contradict the current results, but the key distinction between these studies and our own is that our forced-choice trials used the same words as the single-item trials and paired them based on the single-item responses. This contingency imposed the constraints that defined the obligatory qualitative predictions for both models, and these extra constraints are a key reason why the UVSD model failed here even though it succeeded in past studies.

Our results demonstrate that payoff manipulations can induce strong biases in recognition memory if experimental procedures ensure that participants understand and attend to variations in payoffs (see also Curran et al., 2007; Healy & Kubovy, 1978; Henriques et al., 1994; Snodgrass & Corwin, 1988; Van Zandt, 2000). Target proportion manipulations are a more common method of inducing biases in recognition memory (see the appendix in Broder & Schutz, 2009), but payoff manipulations have the advantage of introducing biases without altering the number of observations across conditions. This advantage was important for our purposes, because having a low trial count for one item type would severely reduce the number of forced-choice trials that could be created. Direct comparisons of proportion and payoff manipulations suggest that they produce similar response patterns (Healy & Kubovy, 1978).

Alternative Model Assumptions

The failed predictions of the 2HT and UVSD models are properly interpreted as evidence against these approaches. We (the authors) were fully prepared to interpret successful predictions as evidence for a given model, and the same logic that supports this impulse motivates the conclusion that failed predictions are evidence against a model. Of course, our results do not provide *definitive* evidence against these approaches or suggest that they no longer serve a useful role in recognition memory research. There are a number of possible “saves” for the models; in other words, one could make alternative assumptions about how our experimental paradigm maps to model mechanisms, potentially negating the model predictions that are out of line with the empirical results. We considered a number of these saves.

First, we considered the possibility that our paradigm encouraged random guessing on the single-item trials. We manipulated response biases by establishing a safe and a risky response on each trial, with safe responses yielding more points for corrects than they deducted for errors, and vice versa for risky responses. To ensure that participants attended to the bias manipulation, we asked them to make two responses on each trial, an initial guess before seeing the test word and a memory-guided response after the test word appeared. As there was no way to avoid frequent errors in the initial guesses, participants needed to consistently select the safe response to avoid losing a lot of points. This enforced guessing procedure might have encouraged participants to persist with a guessing strategy even after the test word appeared. We evaluated this possibility by comparing accuracy and RT data to a similar comparison study that did not use an enforced guessing procedure. Performance levels were very similar across the two studies. We identified an RT cutoff for fast, low-accuracy responses that were potential guesses, and a very low proportion of recognition responses fell below the cutoff, whereas the vast majority of

initial guesses were below the cutoff. Overall, our single-item recognition data showed no signs of excessive guessing after the test word appeared.

Second, we explored whether changes in memory between the single-item and forced-choice trials negated the predicted bias effects on forced-choice accuracy. In our initial implementation of both models, we made the simplifying assumption that memory evidence remained constant between single-item and forced-choice trials based on our practice of interleaving forced-choice trials after every 12 single-item trials. With this setup, an item's appearance in the forced-choice trials comes within seconds of its appearance on the single-item trials with a small number of trials in between, so changes in memory evidence should be minimal between the two. Nevertheless, some changes in memory could plausibly occur based on forgetting across the test list and source confusions (i.e., confusing earlier test presentations with study presentations), so we conducted model simulations that incorporated changes in memory. In short, the model predictions were very robust to fluctuations in memory (Figure 8).

In summary, we found that model predictions were very stable across different assumptions about how our specific task affected model processes. This robustness is not surprising, as the qualitative model predictions are based on selection effects that result from core model mechanisms. For example, the 2HT model predicts lower S-S performance for liberal than conservative bias because a greater willingness to guess “old” means that the pool of recognized targets will include a higher proportion of lucky guesses and a lower proportion of words that were remembered from the list. General forgetting or source confusion processes that introduce changes in memory cannot completely counteract the “contamination” of lucky-guess trials unless there is no correlation between memory states in the single-item and forced-choice trials. Similarly, the UVSD predictions are driven by selection effects “bunching up” the strength

distributions and obscuring differences in mean strength (see Figures 4 and 5), and these selection effects will occur any time words called “studied” (“not studied”) tend to have higher (lower) strength values than the full set of words. Although we cannot rule out all potential “contaminant” processes, the mispredictions of the 2HT and UVSD models appear to be driven by core model mechanisms. In any case, researchers should be searching for models that make accurate predictions for novel experimental tasks, not models that can be saved with *ad hoc* explanations.

Other Models

Our experiment was designed to test the opposite predictions of the 2HT and UVSD models, as outlined in our preregistration document (review link: https://osf.io/kn2tw/?view_only=7c207f36faf745fe85ff47acc75ed648). However, we considered alternative models after results showed that neither of the original models predicted the data. For the Luce (1963) one-low-threshold (1LT) model, one response category always has targets and lures from the same memory state. For example, if participants respond conservatively, then they never respond “studied” for a word that they failed to detect as old, meaning that all lures *and all targets* called “not studied” are in the failed detection state. As a result, the 1LT model predicts that at least one forced-choice trial type, N-N or S-S, must have chance-level performance for a given participant. Results showed that many participants achieved above-chance performance for both trial types, so our results document a failed prediction for the 1LT model.

We also evaluated the dual process signal detection (DPSD) model by challenging it to predict the forced-choice results with parameter values estimated from the single-item data. The DPSD model predicted smaller bias effects on forced-choice accuracy than either the 2HT or UVSD models, so the DPSD predictions were a much better match for the observed data. Of

course this does not negate any previous observations that provide evidence against this model (e.g., Glanzer, Kim, Hilford, & Adams, 1999; Starns, Rotello, & Hautus, 2014; Wixted, 2007), but it does mean that the current results are properly interpreted as support for the DPSD model over the alternative models that we tested. Similarly, O'Connor, Guhl, Cox, & Dobbins (2011) reported evidence for the DPSD model over the UVSD model using a task that combined different types of recognition judgments. In their critical conditions, they presented participants with triplets composed of either two targets and a lure (TTL) or two lures and a target (LLT). Participants completed forced-choice trials in which they were told which type of triplet they would see and they were asked to either select the lure item or the target item for TTL and LLT trials, respectively. Participants also completed oddity-task trials in which they did not know which type of triplet they would get. They were asked to select the item that was from a different stimulus class than the other two, regardless of whether this item was a lure amongst two targets or a target amongst two lures. Results showed that forced-choice performance was comparable for TTL and LLT trials, but oddity task performance was substantially higher for TTL than LLT trials. These results were inconsistent with the UVSD model under a variety of assumed decision rules. In contrast, the DPSD model predicted the trial type differences based on the fact that the TTL trials had two items that could be potentially recollected.

In future work, it will be important to explore how evidence accumulation models (Ratcliff, Smith, Brown, & McKoon, 2016) or discrete state RT models (e.g., Heck & Erdfelder, 2016; Klauer & Kellen, 2018; Starns, 2018) fare with the error-correction recognition paradigm used here. Ideally, this would involve fitting RT distributions for the single-item responses and imposing additional model constraints, perhaps by introducing a speed-accuracy manipulation for the single-item trials. Many RT models have multiple bias mechanisms; for example, the

diffusion model can introduce biases towards “studied” or “not studied” responses with changes in either the starting point of evidence accumulation or the drift criterion that defines the level of memory strength that must be surpassed to provide evidence for a “studied” response (Starns, Ratcliff, & White, 2012). Thus, defining model predictions for bias effects on forced-choice accuracy will require RT data to distinguish alternative bias mechanisms.

Conclusion

This study was inspired by the realization that two popular decision models, the 2HT model and the UVSD model, make strong and opposite qualitative predictions for the effect of response biases in a recently developed “error correction” recognition memory paradigm (Starns et al., 2018). We tested these predictions in two experiments, the second of which was preregistered. Results showed little or no effect of bias on forced-choice accuracy, contradicting the predictions of both models. Although the UVSD and 2HT models are conceptually quite distinct from each other, the ability to distinguish them has often proven to be somewhat challenging (e.g., Broder & Schutz, 2009; Dube & Rotello, 2012; Kellen & Klauer, 2011; Klauer & Kellen, 2010; Starns & Ma, 2018). These challenges were to a large extent due to the continuous reliance on long-established experimental designs and comparison methods (e.g., evaluating the shape of confidence-rating or bias-manipulation ROCs). Overcoming these challenges requires the development of novel designs and new ways of looking at model predictions (e.g., Dube, Starns, Rotello, & Ratcliff, 2012; Kellen & Klauer, 2014, 2015; Starns, Ratcliff, & McKoon, 2012; Starns et al., 2018). More than a “plague in both our houses,” what the present results show is that these new designs pose a number of empirical challenges that will guide the search for better decision models.

References

- Banks, W. P. (1970). Signal detection theory and human memory. *Psychological Bulletin*, 74, 81-99.
- Batchelder, W. H., & Alexander, G. E. (2013). Discrete-state models: Comment on Pazzaglia, Dube, and Rotello (2013). *Psychological Bulletin*, 139, 1204–1212.
- Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review*, 6, 57–86.
- Brysbaert, M., & New, B. (2009). Moving beyond Kucera and Francis: a critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods*, 41, 977-990.
- Bröder, A., & Schütz, J. (2009). Recognition ROCs are curvilinear – or are they? On premature arguments against the two-high-threshold model of recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35, 587-606.
- Busmeyer, J. R., Wang, Y.-M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology*, 44, 171-189.
- Chen, T., Starns, J. J., and Rotello, C. M. (2015). A violation of the conditional independence assumption in the two-high-threshold model of recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41, 1205-1225.
- Curran, T., DeBuse, C., & Leynes, P. A. (2007). Conflict and criterion setting in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 2–17.

- Dube, C., & Rotello, C. M. (2012). Binary ROCs in perception and recognition memory are curved. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38, 130–151.
- Dube, C., Starns, J. J., Rotello, C. M., & Ratcliff, R. (2012). Beyond ROC curvature: Strength effects and response time data support continuous-evidence models of recognition memory. *Journal of Memory and Language*, 67, 389–406.
- Egan, J. P. (1958). Recognition memory and the operating characteristic (Tech. Note No. AFCRC-TN-58–51). Bloomington, IN: Hearing and Communication Laboratory, Indiana University.
- Glanzer, M., Kim, K., Hilford, A., & Adams, J. K. (1999). Slope of the receiver-operating characteristic in recognition memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 25, 500–513.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psycho- physics*. Oxford, England: Wiley.
- Healy, A. F., & Kubovy, M. (1978). The effects of payoffs and prior probabilities on indices of performance and cut-off location in recognition memory. *Memory & Cognition*, 6, 544–553.
- Heck, D. W., Erdfelder, E. (2016). Extending multinomial processing tree models to measure the relative speed of cognitive processes. *Psychonomic Bulletin & Review*, 23, 1440–1465.
- Henriques, J. B., Glowacki, J. M., & Davidson, R. J. (1994). Reward fails to alter response bias in depression. *Journal of Abnormal Psychology*, 103, 460–466.

- Jang, Y., Wixted, J. T., & Huber, D. E. (2009). Testing signal-detection models of yes/no and two-alternative forced-choice recognition memory. *Journal of Experimental Psychology: General*, 138, 291-306.
- Kellen, D., & Klauer, K. C. (2011). Evaluating models of recognition memory using first-and second-choice responses. *Journal of Mathematical Psychology*, 55, 251-266.
- Kellen, D. & Klauer, K. C. (2014). Discrete-state and continuous models of recognition memory: Testing core properties under minimal assumptions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40, 1795–1804.
- Kellen, D. & Klauer, K. C. (2015). Signal detection and threshold modeling of confidence-rating ROCs: A critical test with minimal assumptions. *Psychological Review*, 122, 542–557.
- Kellen, D. & Klauer, K. C. (2018). Elementary signal detection and threshold theory. In E. J. Wagenmakers (Ed.), *Stevens' Handbook of Experimental Psychology and Cognitive neuroscience (4th edition, vol. v)*. New York: Wiley.
- Kellen, D., Erdfelder, E., Malmberg, K. J., Dube, C., & Criss, A. H. (2016). The ignored alternative: An application of Luce's low-threshold model to recognition memory. *Journal of Mathematical Psychology*, 75, 86-95.
- Kellen, D., Singmann, H., Vogt, J., & Klauer, K. C. (2015). Further evidence for discrete-state mediation in recognition memory. *Experimental Psychology*, 62, 40-53.
- Kellen, D., Winiger, S., Dunn, J. C., & Singmann, H. (2019). Testing the Foundations of Signal Detection Theory in Recognition Memory. <https://doi.org/10.31234/osf.io/p5rj9>
- Klauer, K. C., & Kellen, D. (2010). Toward a complete decision model of item and source recognition: A discrete-state approach. *Psychonomic Bulletin & Review*, 17, 465-478.

- Klauer, K. C., & Kellen, D. (2018). RT-MPTs: Process models for response-time distributions based on multinomial processing trees with applications to recognition memory. *Journal of Mathematical Psychology*, 82, 111-130.
- Kroll, N. E., Yonelinas, A. P., Dobbins, I. G., & Frederick, C. M. (2002). Separating sensitivity from response bias: Implications of comparisons of yes-no and forced-choice tests for models and measures of recognition memory. *Journal of Experimental Psychology: General*, 131, 241-254.
- Lockhart, R. S., & Murdock, B. B. (1970). Memory and the theory of signal detection. *Psychological Bulletin*, 74, 100-109.
- Luce, R. D. (1963). A threshold theory for simple detection experiments. *Psychological Review*, 70, 61-79.
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection Theory: A User's Guide*. Lawrence Erlbaum Associates.
- Mosier, C. I. (1951). Problems and designs of cross-validation. *Educational and Psychological Measurement*, 11, 5-11.
- O'Connor, A. R., Guhl, E. N., Cox, J. C., & Dobbins, I. G. (2011). Some memories are odder than others: Judgments of episodic oddity violate known decision rules. *Journal of Memory and Language*, 64, 299-315.
- Parks, C. M., & Yonelinas, A. P. (2009). Evidence for a memory threshold in second-choice recognition memory responses. *Proceedings of the National Academy of Sciences*, 106, 11515-11519.

- Pazzaglia, A. M., Dube, C., & Rotello, C. M. (2013). A critical comparison of discrete-state and continuous models of recognition memory: Implications for recognition and beyond. *Psychological Bulletin, 139*, 1173-1203.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing* (Vol. 124, p. 125).
- Province, J. M., & Rouder, J. N. (2012). Evidence for discrete-state processing in recognition memory. *PNAS Proceedings of the National Academy of Sciences of the United States of America, 109*, 14357-14362.
- Prull, M. W., Dawes, Leslie L. C., Martin, A. M. III, Rosenberg, H. F., Light, L. L. (2006). Recollection and familiarity in recognition memory: Adult age differences and neuropsychological test correlates. *Psychology and Aging, 21*, 107-118.
- Ratcliff, R., Sheu, C-F., & Gronlund, S. (1992). Testing global memory models using ROC curves. *Psychological Review, 99*, 518-535.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences, 20*, 260-281.
- Ratcliff, R., Thapar, A., & McKoon, G. (2004). A diffusion model analysis of the effects of aging on recognition memory. *Journal of Memory and Language, 50*, 408-424.
- Rotello, C. M., Masson, M. E. J., & Verde, M. F. (2008). Type I error rates and power analyses for single-point sensitivity measures. *Perception & Psychophysics, 70*, 389-401.
- Smith, D. G., & Duncan, M. J. J. (2004). Testing Theories of Recognition Memory by Predicting Performance Across Paradigms. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*, 615-625.

- Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, *117*, 34-50.
- Starns, J. J., Dubé, C., & Frelinger, M. E. (2018). The speed of memory errors shows the influence of misleading information: Testing the diffusion model and discrete-state models. *Cognitive Psychology*, *102*, 21-40.
- Starns, J. J., & Ma, Q. (2018). Guessing versus misremembering in recognition: A comparison of continuous, two-high-threshold, and low-threshold models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*, 527-539.
- Starns, J. J., Ratcliff, R., & McKoon, G. (2012). Evaluating the unequal-variance and dual-process explanations of zROC slopes with response time data and the diffusion model. *Cognitive Psychology*, *64*, 1-34.
- Starns, J. J., Ratliff, R., & White, C. N. (2012). Diffusion model drift rates can be influenced by decision processes: An analysis of the strength-based mirror effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*, 1137-1151.
- Starns, J. J., Rotello, C. M., Hautus, M. J. (2014). Recognition memory zROC slopes for items with correct versus incorrect source decisions discriminate the dual process and unequal variance signal detection models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*, 1205-1225.
- Starns, J. J. (2018). Adding a speed-accuracy trade-off to discrete-state models: A comment on Heck and Erdfelder (2016). *Psychonomic Bulletin and Review*, *25*, 2406-2416.

- Starns, J. J. (2021). High- and low-threshold models of the relationship between response time and confidence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 47, 671–684.
- Swets, J. A. (1961). Is there a sensory threshold? *Science*, 134, 168–177.
- Van Zandt, T. (2000). ROC curves and confidence judgments in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 582–600.
- Wixted, J. T. (2007). Dual-Process Theory and Signal-Detection Theory of Recognition Memory. *Psychological Review*, 114, 152-176.

Author Note

Preparation of this article was supported by Grant No. 1454868 from the National Science Foundation. Experiment 2 was preregistered on the Open Science Foundation (OSF; review link: https://osf.io/kn2tw/?view_only=7c207f36faf745fe85ff47acc75ed648)

Table 1

Example block from the memory test illustrating the relationship between the single-item and forced-choice trials

<u>Single-item Trails</u>			<u>Subsequent Forced-Choice Trials</u>	
Word	Stimulus Class	Response	Trial type	Word pair
Compassion	Target	Old	S-S	Restaurant-Way
Restaurant	Lure	Old		
Lady	Lure	New		
Flight	Lure	New		
Obedient	Lure	New	N-N	Lady-Spectacular
Spectacular	Target	New		
Way	Target	Old		
Alien	Target	Old		
Potency	Target	Old	N-N	Disaster-Flight
Disaster	Target	New		

Note: S-S indicates a “Studied”-“Studied” trial and N-N indicates a “Not Studied”-“Not Studied” trial.

Table 2

Mean proportion of “studied” responses on the single-item recognition trials

Item Type	Bias Condition	
	Liberal	Conservative
Target	.76 (.13)	.63 (.18)
Lure	.33 (.21)	.17 (.11)

Note: Values in parentheses are standard deviations.

Table 3

Mean proportion of correct responses on the forced-choice trials

Trial Type	Bias Condition	
	Liberal	Conservative
N-N	.62 (.16)	.63 (.13)
S-S	.76 (.14)	.79 (.15)

Note: Trial types identify forced-choice trials with a target and a lure that were both called “not studied” (N-N) and “studied” (S-S) on the previous single-item trials. Values in parentheses are standard deviations.

Figure Captions

Figure 1. The two-high-threshold (2HT, Panel A) model and Unequal Variance Signal Detection (UVSD, Panel B) model of recognition memory. See text for model explanations.

Figure 2. Schematic of liberal and conservative blocks on the recognition test. Response labels that appear in white and black actually appeared to participants in green and red, respectively.

Figure 3. Predicted forced-choice percent correct for the 2HT and UVSD models across a range of parameter values. In the case of the 2HT, we considered D_o and D_n values ranging from .10 to .90 in steps of .10. In the case of the UVSD, we considered values of μ_t ranging from 0.50 and 3.00 in steps of 0.25 with parameter σ_t fixed to $\sqrt{1.35}$. The plots in Rows 1 and 2 have a line for each value of the memory parameters, and the plots in Row 3 have lines for different combinations of the memory parameters (D_o and D_n for the 2HT model; μ_t , μ_l , and σ_t for the UVSD model).

Figure 4. Example memory strength distributions from the unequal variance signal detection (UVSD) model for all items (A-C) and the corresponding conditional distributions for items called “studied” (D-F). Solid curves show the distributions for lures, dashed curves show the distribution for targets, and the vertical dotted line shows the response criterion. Responding becomes more liberal across panels A-C and D-F. The overlap of the conditional distributions determines forced-choice percent correct on S-S trials. $p(\text{Cor.})$ = predicted forced-choice percent correct.

Figure 5. Example memory strength distributions from the unequal variance signal detection (UVSD) model for all items (A-C) and the corresponding conditional distributions for items called “not studied” (D-F). Solid curves show the distributions for lures, dashed curves

show the distribution for targets, and the vertical dotted line shows the response criterion.

Responding becomes more liberal across panels A-C and D-F. The overlap of the conditional distributions determines forced-choice percent correct on N-N trials. $p(\text{Cor.})$ = predicted forced-choice percent correct.

Figure 6. Percent correct in single-item recognition across a sequence of response time (RT) bins. Each bin had a width of 100 ms and was centered at the position shown on the x -axis. The intervals around each point are 90% credible intervals on a Beta posterior distribution for the proportion estimate. The prior distribution was uninformed (both shape parameters were equal to 1). The dashed vertical line shows the cutoff that we used to identify potential fast guesses. $p(\text{Cor.})$ is the proportion of correct responses (targets called “old” and lures called “new”).

Figure 7. Distributions of predicted bias effects for the two-high-threshold (2HT) and unequal-variance signal detection (UVSD) models with parameters informed by the fits to single-item recognition data. The vertical lines mark the observed bias effects. S-S indicates a “Studied”-“Studied” trial and N-N indicates a “Not Studied”-“Not Studied” trial.

Figure 8. Predicted forced-choice percent correct for different models, across ranges of parameter values. The left column illustrates the predictions of an extended 2HT model, 2HT*, that allows for memory states to differ across single-item and forced-choice judgments. The center column illustrates the predictions of an alternative UVSD model, UVSD*, in which a noise term is added to represent variation in the familiarity of items between the single-item and forced-choice trials. Finally, the right column illustrates the predictions of an extended UVSD model, UVSDCN, in which criterion noise is present. Each line corresponds to the predictions associated to a specific combination of memory parameters. For details on these models, see the main body of text and the supplemental materials.

Figure 9. Estimated forced-choice percent correct on “Studied”-“Studied” (S-S) and “Not Studied”-“Not Studied” (N-N) trials for each participant. Results are from the hierarchical Bayesian model described in the main text. Points show the medians of the posterior distributions for individual participants, and gray lines show 90% credible intervals. The dashed lines mark chance performance.

Figure 10. Predicted forced-choice percent correct for the DPSD model across a range of parameter values. We considered values of $\mu_t = -\mu_t$ ranging from 0.25 and 2.00 in steps of 0.125, and recollection probability R ranging from 0 to .90 in steps .10. Each line corresponds to the predictions obtained for a specific combination of memory parameters, across values of κ .

Figure 11. Distributions of predicted bias effects for the two-high-threshold (2HT) and unequal-variance signal detection (UVSD) models with parameters informed by the fits to single-item recognition data. The vertical lines mark the observed bias effects. S-S indicates a “Studied”-“Studied” trial and N-N indicates a “Not Studied”-“Not Studied” trial.

Figure 1

Figure 2

Figure 3

Figure 4

Figure 5

Figure 6

Figure 7

Figure 8

Figure 9

Figure 10

Figure 11

Appendix: Significance Tests

Our Bayesian analyses indicated that we observed large bias effects in single-item recognition but null or negligible bias effects on forced-choice accuracy. Here, we evaluate the bias effects with more traditional significance tests and confidence intervals, using $\alpha = .05$. t tests showed significant bias effects for both hit rate, $t(146) = 8.63, p < .001, CI [0.09, 0.15]$, and false-alarm rate, $t(146) = 8.75, p < .001, CI [0.11, 0.17]$. Three participants could not be included in the analyses of forced-choice accuracy because they had no forced-choice trials in one of the conditions (e.g., they never called a target “new” with liberal bias and thus had no N-N trials for liberal blocks). As in the main analyses, we defined bias effects as forced-choice percent correct in the liberal condition minus the conservative condition. The bias effect was nonsignificant for N-N trials, $t(144) = -1.30, p = .19, CI [-0.01, 0.05]$. For S-S trials, there was a significant bias effect, $t(144) = -2.18, p = .03$, but the confidence interval included functionally null effect sizes and ruled out anything larger than a small effect $[0.003, 0.05]$. Comparing the relative effect size on S-S and N-N trials produced a nonsignificant result, $t(144) = 0.43, p = .67, CI [-0.03, 0.05]$.