Does source memory exist for unrecognized items?

Julian Fox
University of Melbourne

Adam F. Osth
University of Melbourne

Address correspondence to:

Julian Fox (E-mail: jfox1@student.unimelb.edu.au)

Author Note

Data and model code from this article can be found on our Open Science Framework (OSF) page (<https://osf.io/my93p/>). We would like to thank Simone Malejka and Jeffrey Starns for generously providing their data. This work was supported by a grant from the Australian Research Council, ARC DE170100106, awarded to Adam Osth.

Abstract

In episodic memory research, there is a debate concerning whether decision-making in item recognition and source memory is better explained by models that assume all-or-none retrieval processes or continuous underlying strengths. One aspect in which these classes of models tend to differ is their predictions regarding the ability to retrieve contextual details (or *source* details) of an experienced event, given that the event itself is not recognized. All-or-none or high-threshold models predict that when items are unrecognized, source retrieval is not possible and only guess responses can be elicited. In contrast, models assuming continuous strengths predict that it is possible to retrieve the source of unrecognized items, albeit with low accuracy. Empirically, there have been numerous studies reporting either chance accuracy or above-chance accuracy for source memory in the absence of recognition. Crucially, studies presenting recognition and source judgements for the same item in immediate succession (*simultaneous design*) have revealed chance-level accuracy, while studies presenting a block of recognition judgements followed by a block of source judgements (*blocked design*) have revealed slightly above-chance accuracy. Across three sets of experiments involving multiple design manipulations, the present investigation demonstrated: (a) that source memory for unrecognized items is indeed higher in blocked designs; (b) that evidence for the effect in blocked designs is likely artifactual due to item memory changing between blocks; and (c) that the effect does exist in simultaneous designs, but is highly subtle and is more easily detected when uncertainty in the participant-level data is low or is accounted for in a hierarchical Bayesian model. It is suggested that findings of a null effect in the prior literature may be attributable to design elements that hinder source memory as a whole, and to high degrees of uncertainty in the participant-level source data when conditioned on unrecognized items.

Keywords: recognition memory, source memory, high-threshold models, signal detection

models, hierarchical Bayesian modelling

Does source memory exist for unrecognized items?

In episodic memory research, there is a distinction between memory for learned information and memory for the context in which information was learned (Johnson et al., 1993). These two types of memory are investigated in the laboratory using the item recognition and source memory paradigms respectively. In the item recognition paradigm, participants study a list of items and at test are required to discriminate between items that they have just studied (i.e., targets), and items that they did not study (i.e., lures). In the source memory paradigm, participants study a list of items that are presented in different source contexts (e.g., different font colours, different locations, different gendered speakers). At test, participants are presented with the same items, but removed from their source context, and are required to identify the source in which each item was initially presented.

Many computational models have been developed to jointly account for item recognition and source memory (e.g., Banks, 2000; Batchelder & Riefer, 1990; DeCarlo, 2003; Glanzer et al., 2004; Hautus et al., 2008; Hilford et al., 2002; Klauer & Kellen, 2010; Osth et al., 2018; Slotnick & Dodson, 2005; Yonelinas, 1999). While a number of these models have been successful in accounting for data trends across various tasks and conditions, particularly the shapes of receiver operating characteristics (ROCs), they each make different assumptions about the mechanisms that underly item and source memory. Some of the models describe item and source decision-making in terms of all-or-none retrieval processes or high-thresholds (Batchelder & Riefer, 1990; Klauer & Kellen, 2010), other models posit that item and source decisions are based on continuous memory strengths (Banks, 2000; DeCarlo, 2003; Glanzer et al., 2004; Hautus et al., 2008; Hilford et al., 2002; Osth et al., 2018), and some involve a combination of both high-thresholds and continuous strengths (Yonelinas, 1999).

Although these models jointly predict many phenomena in item and source memory, the

critical distinction between them is whether they posit that it is possible to retrieve the source context of an item that is not recognized (Starns et al., 2008). In other words, whether participants are capable of identifying the sources of items that they have labelled as “NEW” in a recognition test with above-chance accuracy. This phenomenon has been investigated in numerous publications, with results sometimes demonstrating source memory for unrecognized items (e.g., Starns et al., 2008) or a complete lack thereof (e.g., Malejka & Bröder, 2016). As we will describe, there are numerous design and analysis decisions that may have influenced this discrepancy.

Model Predictions

An inherent feature of all-or-none or high-threshold models is that it is only possible to perform source retrieval for recognized items. A demonstration of this can be found in the multinomial processing-tree models of Batchelder and Riefer (1990). Multinomial models represent item and source memory in terms of underlying cognitive states that are experienced discretely. Each cognitive state has a certain probability of being experienced given the nature of the presented item. Furthermore, the possible cognitive states are connected in a branching tree-like structure, such that the experience of a particular state is conditional on the experience of the preceding state in the tree. Figure 1 illustrates the structure of a typical multinomial model for an item/source test, given that an item from “Source A” is presented. The leftmost label represents the presented item, and the rightmost labels indicate the source judgements that can be made for the item. As can be seen, when a Source A item is presented, there is a certain probability that the item will be recognized (D_A), and a certain probability that it won't be ($1 - D_A$). If the item is unrecognized, then there is a certain probability that the item will be assumed to be “OLD” (b), and a certain probability that the item will be assumed to be “NEW” ($1 - b$). If the unrecognized item is assumed to be “OLD”, then there is certain probability that the participant will correctly

guess the item is from “Source A” (g), and a certain probability that they will guess “Source B” ($1 - g$), before responding accordingly. Given that the only two cognitive states that can result from these unrecognized items are guesses, multinomial models explicitly make the assumption that is not possible to perform source retrieval for unrecognized items.

Figure 1

Representation of a Multinomial Model for Item/Source Memory

Note. The leftmost label indicates that the test cue is a “Source A” item; the rightmost labels indicate the possible responses that a participant can make for the item in the source test. The processing tree structure in-between illustrates the succession of cognitive states that lead to each response. D_A = probability of recognizing the Source A item; d_A = probability of recognizing the item’s source given it is recognized; b = probability of assuming the item is “OLD” given it is not recognized; a = probability of correctly guessing the source of the item given it is recognized; g = probability of correctly guessing the source of the item given it is not recognized.

The dual-process model of Yonelinas (1999) differs from the multinomial models described above as it incorporates both high-thresholds and continuous strengths, yet is comparable in that it also predicts no source memory for unrecognized items. The model assumes that two distinct cognitive processes underly memory judgements: familiarity and recollection. It is assumed that source judgements rely primarily on recollection processes, and that these are governed by a high-threshold mechanism. As with multinomial models, the

all-or-none nature of the high-threshold mechanism is such that information relevant to the study event can be successfully retrieved if the threshold is surpassed, but that no recollective information is available below the threshold and only guess responses can be made. Given that source memory is driven purely by recollection, if an item is not recognized (i.e., it is not familiar), then the threshold will not be surpassed and there will be no source information available, resulting in a guess response. One exception to this is if the model allows source information to be encoded as a feature of a studied item – such that the item and source information are “unitized” – and not just as an arbitrary contextual association (Diana et al., 2008). In this case, the familiarity component of the dual-process model would also contribute to source memory performance. Given that the familiarity process is governed by continuous strengths, the incorporation of such a *unitized familiarity* mechanism would predict above-chance source accuracy for unrecognized items.

One class of models that does predict above-chance source accuracy for unrecognized items are those incorporating multivariate signal detection theory (SDT; Banks, 2000; DeCarlo, 2003; Glanzer et al., 2004; Hilford et al., 2002; Starns et al., 2008). Unlike multinomial models (Batchelder & Riefer, 1990) and the dual-process model (Yonelinas, 1999), multivariate SDT models do not involve high-thresholds or discrete states in general, rather they assume that item and source decisions are based on continuous memory strengths. One such model, the two-dimensional signal detection model (2D-SDT; Hilford et al., 2002) is depicted in Figure 2. The model is comprised of three bivariate normal distributions, representing Source A items, source B items, and “NEW” items (labelled A, B, and N respectively). The three distributions are arranged in a triangular formation on a Cartesian plane with the x-axis representing recognition evidence (i.e., evidence in favour of “OLD” items, either Source A or Source B), and the y-axis representing source evidence (i.e., evidence in favour of Source B items). While the N distribution occupies a low position on

the recognition dimension, the A and B distributions each occupy a high position.

Furthermore, while the B distribution occupies a high position on the source dimension, the A distribution occupies a low position, and the N distribution lies in-between. Note that the A and B distributions are elliptical, reflecting the observation that item and source strengths are positively correlated (Glanzer et al., 2004). Each item presented at test – be it from distribution A, B, or N – can be represented by a point on the plane that corresponds to its degree of recognition and source evidence. In order to make a recognition judgement for a test item, a recognition criterion (c_R) is established at a point along the recognition dimension; items that fall to the left or right of the criterion are deemed “NEW” or “OLD” respectively. Likewise, in order to make a source judgement for a test item, a source criterion (c_S) is introduced along the source dimension; items that fall to the left or right of the criterion are judged to originate from Source B or A respectively.

Figure 2

The Two-Dimensional Signal Detection Model for Item and Source Memory (2D-SDT; Hilford et al., 2002)

Note. The three ovals represent evidence distributions for each type of test item (“NEW” items, Source A items, and Source B items). The vertical lines represent recognition criteria (liberal and conservative respectively) and the horizontal lines represents the source criterion.

When an item is unrecognized at test (i.e., it is deemed “NEW”), the ability to correctly identify its source is reflected in the overlap between the A and B distributions below the recognition criterion, with greater overlap indicating lower source discriminability. Due to the correlation between evidence values on the item and source dimensions, the overlap between the two source distributions increases as recognition evidence decreases (see Figure 2). If the recognition criterion is highly liberal, then only items with very low recognition evidence will be judged as “NEW”, and consequently there will be virtually no source memory for unrecognized items. However, if a more conservative criterion is established, then there will be some items with higher recognition evidence judged as “NEW”, thus allowing some source accuracy for unrecognized items¹ (Starns et al., 2008).

While multivariate SDT models by nature allow for the possibility of source memory for unrecognized items, Hautus et al. (2008) investigated several variants of their bivariate signal detection model and ironically found the most successful one included an ad hoc guessing mechanism for source memory when items were unrecognized. They argued that when participants complete a dual-rating task that involves a recognition judgement and source judgement for each item (Yonelinas, 1999), it is unlikely that any “NEW” judgement ascribed to an item would lead to an effortful source judgement for that item, as participants would consider the judgement meaningless. Accordingly, they observed superior model fit when their model incorporated the assumption that “NEW” judgements for items in the

¹ Although not depicted in Figure 2 for simplicity’s sake, the A and B distributions of the 2D-SDT model should theoretically extend beyond the point where they overlap, thus forming a roughly ‘X’ shape. A counterintuitive (and likely implausible) implication of this is that with an extremely liberal recognition criterion, participants should make systematically false source judgements for unrecognized items (i.e., systematically attribute source A items to source B and vice versa).

recognition stage would lead to guess responses in the source stage, rather than decisions based upon continuous latent strengths. Although the latter model allowed source accuracy to decrease as item confidence decreased, it was unable to capture the very abrupt decrease in source performance that resulted when an item was unrecognized. By allowing the assumption that unrecognized items elicit guesses, their model was able to capture this pattern.

Just as the model of Hautus et al. (2008) stands as an exception to the generalization that continuous strength models predict source memory for unrecognized items, it is also important to acknowledge that some threshold models could be capable of predicting source memory in the absence of recognition. While high-threshold models postulate that retrieval is an all-or-none process, with retrieval failures resulting in guess responses, low-threshold models (Kellen et al., 2016; Luce, 1963) postulate that retrieval is a noisy process that can produce either a high or low state of evidence (Starns & Ma, 2018). If a low-threshold source memory model were devised, then it would be possible for a target to produce sub-threshold item evidence (leading to a “NEW” response), but for the target to produce above-threshold source evidence (leading to an accurate source response).

Empirical Findings

While the majority of source memory models possess the inherent feature or incorporate the post-hoc assumption that source retrieval cannot be performed for unrecognized items, the empirical findings regarding this ability are quite mixed. A number of studies have observed chance-level source accuracy for unrecognized items (Bell et al., 2017; Malejka & Bröder, 2016; Malejka et al., 2018; Onyper et al., 2010; Slotnick & Dodson, 2005; Yonelinas, 1999), while others have observed above-chance accuracy for such items (Ceci et al., 2010; Chen et al., 2018; Osth et al., 2018; Starns et al., 2008; Starns et al., 2013; Starns & Ksander, 2016). In an attempt to make sense of this discrepancy, the following sections will provide an

overview of the methodological details and empirical results of each of the aforementioned studies. A summary of these observations, including methodological details of each study (e.g., response modalities and source manipulations), is provided in Table 1.

Table 1

Summary of Experiments Observing Chance or Above-Chance Source Accuracy for Unrecognized Items, Including Methodological Details on Design Type, Rating Scale, Source Manipulation, Method of Analysis, and Other Conditions

	Design	Ratings	Source	Analysis	Misc.
Chance Accuracy					
Yonelinas (1999, Exp. 2)	S	C	List	$d' *$	-
Slotnick and Dodson (2005, Exp. 1)	S	C	Speaker	ROC	-
Slotnick and Dodson (2005, Exp. 2)	S	C	Speaker	ROC	-
Onyper et al. (2010, Exp. 2)	S	C	Location	ROC	-
Malejka and Bröder (2016)	S	Y	Task	P(c)	-
Malejka et al. (2018)	S	Y	Location	P(c)	Short/Long Recognition
Bell et al. (2017, Exp. 1)	S	Y	Colour	Model	-
Bell et al. (2017, Exp. 2)	S	Y	Colour	Model	-
Bell et al. (2017, Exp. 3)	S	Y	Label	Model	-
Zhou et al. (2021)	B	CO	Location	Rayleigh	-
Above-Chance Accuracy					
Starns et al. (2008, Exp. 1)	B	Y	Size	P(c)	Conservative Bias Only
Starns et al. (2008, Exp. 2)	B	Y	Task	P(c)	Conservative Bias Only
Ceci et al. (2010, Exp. 1)	B	Y	Location	P(c)	-
Ceci et al. (2010, Exp. 2)	RB	Y	Location	P(c)	-
Starns et al. (2013,	B	C	Speaker	$d' *$	Strong Items Only

Exp. 2a) Starns et al. (2013, Exp. 2b)	B	C	Speaker	d' *	Strong Items Only
Osth et al. (2018, Exp. 3)	B	C	Location	Model	-
Chen et al. (2018, Exp. 1)	B/RB	Y	Speaker	P(c)	Context-Consistent Targets
Chen et al. (2018, Exp. 2)	J	Y	Speaker	P(c)	Context-Consistent Targets

Note. Starns et al. (2008, Exp. 3) and Starns and Ksander (2016, Exp. 2) excluded for not providing source judgements for lures rated “NEW” in a simultaneous design (i.e., they employed the simultaneous-feedback design). Cook et al. (2006), Kurilla and Westerman (2010), and Brainerd et al. (2012) excluded for not conditioning source accuracy on unrecognized items. S = simultaneous design; B = blocked design; RB = reverse blocked design; J = jumbled design; C = confidence ratings; Y = yes-no ratings; CO = continuous-outcome task; P(c) = percent/probability correct.

*Analyses were conducted independently (i.e., not in the original article) using d' .

Chance Accuracy. One of the first experiments to test participants' recognition and source memory ability for the same items was Experiment 2 of Yonelinas (1999). The experiment utilized a *simultaneous* design to testing item and source memory, whereby participants provide an initial recognition judgement and a subsequent source judgement for each item before the next item is presented (see Figure 3). The experiment also involved participants providing recognition and source ratings via a confidence scale. Although source memory for unrecognized items was not analyzed in the original article, Slotnick and Dodson (2005) re-analyzed these data and revealed that source memory discriminability was only above chance for items recognized with moderate to high confidence. Subsequently, similar experiments involving a simultaneous design and confidence ratings have observed linear source memory ROCs when unrecognized items are included (indicating that accuracy is at chance; Slotnick & Dodson, 2005), and linear source memory ROCs for both unrecognized and moderately recognized items on their own (Onyper et al., 2010).

Figure 3

Schematics of Three Design Types That Can Be Used to Investigate Item and Source Memory for the Same Items

Simultaneous Design

Blocked Design

Reverse Blocked Design

Note. Each rectangle represents a single test trial, with blue indicating a recognition trial and red indicating a source trial. The order in which trials are presented to participants is shown from left to right.

The finding that source accuracy is at chance for unrecognized items has been replicated by a few recent studies using the simultaneous design, but with each requiring participants to respond using binary *yes-no* judgements instead of confidence scales (Bell et al., 2017; Malejka & Bröder, 2016; Malejka et al., 2018). Malejka and Bröder (2016) demonstrated a complete lack of source memory for unrecognized items across both liberal and conservative conditions of response bias, thus contradicting the predictions of the 2D-SDT model. Malejka et al. (2018) acknowledged that this finding may have been confounded by participants providing automatic guess responses to items they had rated “NEW”, given their knowledge that “NEW” items do not have a source. To ensure that participants were motivated to perform source retrieval for each item, they replaced lures with additional targets that were studied under imperceptibly short presentation times, but nonetheless had a source. Despite

the increased motivation to respond effortfully to items rated “NEW”, participant still demonstrated a lack of source memory for unrecognized items. Bell et al. (2017) also observed a lack of source memory for unrecognized items (as measured by a multinomial source monitoring model), even when the studied source information was emotionally valent.

Above-Chance Accuracy. In terms of experiments that have revealed above-chance source accuracy for unrecognized items, a small number have utilized the aforementioned simultaneous design (Starns et al., 2008, Exp. 3; Starns & Ksander, 2016). Note that these studies did not report their own analyses of source memory for unrecognized items; rather we established these findings via independent analyses of their data (see Table 1). However, many other experiments revealing above-chance source accuracy for unrecognized items have utilized a *blocked* design (Ceci et al., 2010, Exp. 1; Chen et al., 2018; Osth et al., 2018; Starns et al., 2008, Exps. 1 & 2; Starns et al., 2013) that involves all items being tested for recognition, before a source memory test even occurs, followed by all items being tested for source memory (see Figure 3). Ceci et al. (2010, Exp. 2) and Chen et al. (2018, Exp. 1) also observed above-chance source accuracy for unrecognized items using a *reverse blocked* design (i.e., a block of source testing followed by a block of recognition testing). Furthermore, Chen et al. (2018, Exp. 2) observed the phenomenon with a unique *jumbled* design that involved a continuous stream of recognition and source judgements, with recognition and source judgements for the same item being separated by six or more trials (recognition judgements would precede source judgements for half the items, and the reverse order would occur for the other half). Of the studies mentioned above, some utilized confidence judgements (Osth et al., 2018; Starns et al., 2013; Starns & Ksander, 2016), while others utilized yes-no ratings (Ceci et al., 2010; Chen et al., 2018; Starns et al., 2008).

It is also important to highlight that certain of these experiments have observed above-chance source accuracy for unrecognized items only under certain conditions. For instance,

Starns et al. (2008) observed an effect when participants responded conservatively, Ceci et al. (2010) observed a greater effect for young children, the Starns et al. (2013) data revealed an effect for items that were strengthened (i.e., repeated during study), and Chen et al. (2018) observed an effect when targets were *context-consistent* (i.e., they shared their source with other items from the same semantic category).

One experiment that should be highlighted is Experiment 3 of Osth et al. (2018). While a Bayesian statistical analysis of their data was agnostic as to whether source discriminability (d') for unrecognized items was above chance ($M = 0.15$, $BF_{10} = 1.13$), a model selection procedure involving various hierarchical Bayesian SDT models indicated strong evidence for a model where d' was not zero. When fit to the data, this model measured performance to be slightly above chance for both high frequency ($M = 0.10$, 95% highest density interval, or HDI: [.02, .17]) and low frequency ($M = 0.21$, 95% HDI: [.09, .33]) words. They suggested that the discrepancy between the two analyses was likely due to the fact that the former was unable to account for the hierarchical structure of the data (i.e., with trials nested within participants, and participants nested within experimental groups), while the latter could. They argued that accounting for the hierarchy of data was pertinent to their research as recognition performance was quite high within the sample, meaning that there were few unrecognized items to analyse per participant, and consequently, a high degree of uncertainty in the participant-level estimates of source accuracy for those items. Given that traditional analysis methods only use point estimates to represent each participant, they would be unable to account for this participant-level uncertainty, and would instead treat each participant as equal when calculating group-level performance. In contrast, hierarchical Bayesian models quantify participant-level uncertainty and allow participants with greater uncertainty in their parameter estimates to be devalued when estimating group-level parameters.

Furthermore, hierarchical Bayesian models (as opposed to standard Bayesian models)

involve a process of *shrinkage*, whereby uncertain estimates at the participant level (particularly outliers) are adjusted so that they become closer to the overall group estimates (Boehm et al., 2018). As Boehm et al. (2018) observed, accounting for the hierarchical structure of data is non-trivial, as it often has a direct impact on the conclusions of research. They showed that modelling strategies (either Bayesian or frequentist) that ignore the hierarchy of data will regularly overestimate group-level variance and consequently be biased towards the null hypothesis, as well as suffer from an overconfidence bias. Furthermore, these biases are most severe when datasets include many participants, but few trials per each.

A potential issue, however, with some of the aforementioned studies is that participants were required to make source judgements for targets that they judged to be “NEW”, but were not required to make source judgements for lure items that they deemed “NEW” (Starns et al., 2008, Exp. 3; Starns et al., 2013; Starns & Ksander, 2016), or in some cases, for lure items at all (Ceci et al., 2010; Osth et al., 2018; Starns et al., 2008, Exp 1 & 2). Malejka and Bröder (2016) highlighted this as a potential confound since any item deemed “NEW” that leads to a prompt for a source judgement would provide implicit feedback that the item was indeed a target (see also Kellen et al., 2014). By informing participants that a previously “unrecognized” item is a target, participants may attempt a second retrieval of the item and consequently re-evaluate the item as “OLD” before attempting source retrieval. In fact, Malejka and Bröder (2016) were able to replicate the findings of Starns et al. (2008, Experiment 3) when lure items judged as “NEW” did not require a source judgement (i.e., when using a simultaneous-feedback design), but were unable to do so when all items required a source judgement (i.e., when using a standard simultaneous design with no implicit feedback). That said, it is likely that this potential confound would only have an effect on simultaneous designs (i.e., the simultaneous-feedback design), as the large gap between recognition and source judgements in a blocked design would attenuate any implicit feedback

about the true status of an item. It should also be noted that while including source judgements for lures rated “NEW” in a simultaneous design is necessary to avoid the possibility of second retrieval attempts, doing so may introduce a separate confound in that participants may become less motivated to perform source retrieval for items they rate “NEW”.

Other Findings. Another study that helps inform the debate is Cook, Marsh and Hicks (2006). Across four experiments, participants were required to recall a target word after being presented with an associated cue word (i.e., cued recall), and to then identify the source context of the target word regardless of whether it was correctly recalled (c.f., simultaneous designs). If cued recall was unsuccessful, then participants were subsequently required to recognize the target word using four-alternative forced choice (4AFC), and then attempt source recognition for a second time. Across the four experiments, it was found that participants could reliably identify the source of an unrecalled target if the target was subsequently recognized. While source identification of the unrecalled targets prior to 4AFC recognition was generally at chance, above-chance accuracy was observed in Experiments 2 and 4 when participants were presented with strengthened targets and/or completed the task under full disclosure. While it is difficult to compare these findings to those mentioned above, by virtue of the fact that source accuracy was not explicitly conditioned on items receiving a “NEW” judgement, the study demonstrates that source memory can exist in the absence of recognition memory.

Also relevant are two experiments by Kurilla and Westerman (2010) that required participants to complete fragments of words that may or may not have appeared on the study list. Following each fragment completion prompt, participants provided a recognition judgement on a confidence scale, and if they said “OLD”, an additional source judgement on a binary scale (i.e., a simultaneous design). While participants exhibited above-chance source

accuracy for word fragments that were not completed but later recognized, no source judgements were required for unrecognized items and so the results cannot be directly compared to others. Additionally, while Brainerd et al. (2012) did not provide recognition and source tests for the same items across their experiments, they observed inferior fits using the conjoint process dissociation (CPD) model when the assumption that source retrieval does not require item memory was removed from the model.

The Impact of Blocked and Simultaneous Designs

Given the wealth of research findings suggesting that source memory is both possible and not possible in the absence of recognition memory, it may appear that the dilemma is far from being resolved. However, inspection of Table 1 reveals that there is one design element that is consistent among experiments reporting chance accuracy, and an opposing design element that is consistent among experiments reporting the above-chance accuracy. If one excludes experiments that used the simultaneous-feedback design (i.e., where source judgements were not required for lures rated “NEW”; Starns et al., 2008, Exp. 3; Starns & Ksander, 2016, Exp. 2) or did not condition source accuracy on unrecognized items (Brainerd et al., 2012; Cook et al., 2006; Kurilla & Westerman, 2010), then it can be seen that almost all experiments reporting chance source accuracy for unrecognized items have utilized a standard simultaneous design (i.e., with no implicit feedback), and almost all experiments reporting above-chance accuracy for unrecognized items have utilized a blocked design.

Why would it be that simultaneous designs lead to chance-level source accuracy for unrecognized items, but that blocked designs lead to above-chance accuracy? Barring the possibility that this observed pattern is merely coincidental, there are at least two plausible explanations for its occurrence. One potential explanation is that the use of a simultaneous design may demotivate participants from responding effortfully in certain source judgements. Consider that in a simultaneous design, participants make a recognition judgement for an

item and are then immediately required to make a source judgement for the same item. If a participant identifies an item as a lure and are then immediately required to identify the source of the item, they may consider the judgement meaningless (as lures do not have a source), and simply resort to a guess response (Hautus et al., 2008). Even if lures are substituted with briefly presented items that do have a source (see Malejka et al., 2018), the very act of providing a negative recognition judgement may alleviate participants' sense of responsibility to attempt effortful source retrieval. While blocked designs also require participants to make source judgements for items that they have rated as “NEW”, the large gap between these two judgements will mean that when a participant is required to make a source judgement their initial recognition response will not be readily available. Note that this same feature of blocked designs applies to the jumbled design of Chen et al. (2018, Exp. 2), and may help account for their finding of source memory for unrecognized items, despite the design's surface-level similarity to simultaneous designs. In a nutshell, it could be argued that source memory exists for unrecognized items, but that traditional simultaneous designs provoke enough guess responses from participants to mask any observable accuracy.

While flaws in the simultaneous design may be responsible for disguising source accuracy for unrecognized items, another convincing explanation for the empirical distinction between blocked and simultaneous designs is that considerable changes in item memory and/or participant behaviour may occur between the two blocks in blocked designs. While analyses of source memory for unrecognized items are dependent on the assumption that item memory and participant behaviour are consistent at the times of item and source retrieval, there is a multitude of evidence to suggest that this would not be the case in blocked designs. For instance, research suggests that the context cues used for retrieval continually shift throughout periods of testing as additional test items are retrieved (Estes, 1955; Howard & Kahana, 2002). Context cues are even likely to change between recognition and source

blocks due to differences in the instructions and composition of the test lists (lures only appear in recognition tests). The implication of this is that an item may be recognized in one block with a particular set of context cues, but not be recognized in another block when different context cues are used. Throughout testing, the exposure to more and more test items will also lead to a gradual accumulation of item interference that impacts the accuracy of recognition judgments (Gillund & Shiffrin, 1984). It is therefore possible that an item could be recognized in one block, but be obscured from memory in the subsequent block due to an influx of interference from the intervening items. On top of that, it is also likely that the probability of remembering a target item is higher in the second block than the first due to the fact that each target would have been learned twice by the second block (i.e., during study and during the first block).

In terms of changes in participant behaviour, there is evidence to suggest that intraindividual fluctuations in response criteria can occur throughout an experiment (Benjamin et al., 2009). It is possible that an item with moderate familiarity could be recognized as “OLD” in one block, but then judged as “NEW” in another block if a particularly conservative criterion were adopted. Furthermore, it has been shown that fluctuations in attention and motivation can occur throughout an experiment (deBettencourt et al., 2018). If participants responded carelessly to a few trials in one block, but responded with proper attention to the same items in the other block, then it would be likely for some recognized items to be mislabelled as unrecognized in the alternate block or vice versa. If any of the above changes (i.e., changes in context cues, item interference, response criteria, and attention/motivation) occurred in such a way that an item was judged as “NEW” during a block of recognition testing, but then implicitly judged as “OLD” prior to source retrieval in a block of source testing, then it could produce artifactual source accuracy for unrecognized items. In other words, it may not be possible to perform source retrieval in the absence of

item memory, but frequent changes in memory and/or responding in blocked designs may disguise regular source memory for recognized items as source memory for unrecognized items.

The Present Experiments

Verification of Findings Between Designs. The first two experiments in the present investigation aim to verify the notion that the finding of source memory for unrecognized items is primarily dependent on experimental design. Specifically, we aim to replicate the presence of the effect while employing a blocked or reverse blocked design, as well as the absence of an effect while employing a simultaneous design. Each experiment will involve a within-subjects manipulation of design, such that all participants complete blocks of recognition and source testing for the same study items under each of the three conditions (see Figure 3). The benefit of using a within-subjects design is that it can be revealed whether the contrasting results found in prior experiments are due to sample noise and/or extraneous methodological details, or due to selection of a blocked/simultaneous design. Note that this would be, to our knowledge, the first study to manipulate blocked and simultaneous designs within-subjects. Each experiment will also involve a manipulation of word frequency. The purpose of this manipulation is to verify prior findings of a low frequency word advantage for source d' for unrecognized items (Osth et al., 2018) and for source memory more generally (Glanzer et al., 2004; Mulligan & Osborn, 2009), as well as to provide extra constraint when our data are fit with computational models.

While the two experiments will be largely identical, the first experiment (Exp. 1a) will require participants to respond using six-point confidence scales, and the second experiment (Exp. 1b) will require participants to respond using binary yes-no judgements (i.e., “OLD” or “NEW” for item recognition, and “Source A” or “Source B” for source memory). While there does not appear to be any systematic differences in findings between studies using confidence

ratings and yes-no ratings (see Table 1), it is noteworthy that no study has used both methods, and therefore it remains unknown whether response methods impact analyses of source memory for unrecognized items. Such a notion is entirely plausible as it been shown that the cognitive burden of translating subjective evidence to ratings on lengthy response scales can introduce *decision noise* and consequently reduce recognition performance (Benjamin et al., 2013). If the key differences between blocked and simultaneous designs can be replicated across both response modalities, then that would provide further credence to notion that source memory for unrecognized items is primarily dependent on experimental design.

In addition to analysing the data from these experiments using statistical methods, we analyzed our data using hierarchical Bayesian SDT models. As mentioned earlier, hierarchical Bayesian models are unique in that they estimate parameters at both the group level and participant level, and allow uncertain participant-level parameters to be adjusted so that they are closer to the group-level parameters (Boehm et al., 2018; Wagenmakers, 2009). This is beneficial as it is common for the amount of unrecognized items to vary greatly by participant, and consequently, for participant-level estimates of source memory for unrecognized items to be highly uncertain. This phenomenon is non-trivial as it has been shown that hierarchical Bayesian estimates of source memory for unrecognized items can differ to what is observable in an analysis of point estimates (Osth et al., 2018). It may be the case that the true impact of experimental design on source performance cannot be determined via an analysis of point estimates, but can be determined via a hierarchical Bayesian analysis.

Validation of Blocked Designs. To preface the results of Experiments 1a and 1b, evidence of source memory for unrecognized items was established in each of the three design types. However, the strongest and most consistent evidence was established in the two blocked designs. This was a concern as any analysis of source memory for unrecognized items is dependent on the assumption that the memory state during item retrieval matches the

memory state during source retrieval, but the nature of experimental blocking may mean that memory during both these phases is different. For instance, participants may not recognize an item during the first block, but then implicitly recognize it on another attempt during the second block. Conversely, participants may implicitly recognize an item at source retrieval during the first block, but then due to item interference, fail to recognize the item during the second block. If this were the case, then the evidence established in the blocked designs may have been largely artifactual.

Following this concern, the second set of experiments of the present investigation seek to determine whether the findings for blocked and reverse blocked designs are legitimate or artifactual. In doing so, Experiments 2a and 2b (utilizing confidence and yes-no ratings respectively) will involve participants' memory for study items being explicitly measured at both blocks in blocked and reverse blocked designs. Instead of providing source memory judgements during one of the two blocks, participants will indicate whether they remember or forget seeing each target item on the study list. This task is distinct from recognition in that no lures are presented, but will still provide a good indication of whether participants' memory for items changes from one block to the next (Koop, Criss & Malmberg, 2014). If it turns out that a high proportion of unrecognized targets are “remembered” in the alternate block, and that this occurs often enough to generate the levels of source memory for unrecognized items that were observed in the blocked conditions of Experiments 1a and 1b, then that would suggest that source memory for unrecognized items in blocked designs is actually source memory for recognized items that happened to be unrecognized at some other point.

Replication of Findings From the Simultaneous Design. To preface the findings from Experiments 2a and 2b, it was revealed that memory for study items frequently changed from one block to the next, thus diminishing the validity of blocked designs in exploring source

memory for unrecognized items. Given this knowledge, Experiments 3a and 3b (utilizing confidence and yes-no ratings respectively) were devised to establish highly reliable and valid estimates of source memory for unrecognized items from the simultaneous design (i.e., the design that we still have faith in). To do so, we collected three times the amount of data per participant for the simultaneous design than in the first set of experiments, thus establishing much more robust participant-level estimates of source memory for unrecognized items. Furthermore, we collected roughly three times the amount of participants per experiment than in Experiments 1a and 1b, thus establishing much more robust estimates of the effect at the group-level. As with the first experiments, the data of Experiments 3a and 3b will be investigated with both statistical methods and hierarchical Bayesian methods.

The data from each of the three sets of experiments can be found on our Open Science Framework (OSF) page (<https://osf.io/my93p/>).

Experiments 1a and 1b

Method

Participants. Both experiments were approved by the Human Research Ethics Committee of The University of Melbourne (Ethics ID no. 12033). For Experiment 1a, the participants were $N = 58$ students aged 18 or older at the University of Melbourne. For Experiment 1b, the participants were $N = 49$ students from the same population, none of which had partaken in Experiment 1a. All participants were recruited through the Melbourne School of Psychological Sciences Research Experience Program (REP). The participants partially fulfilled their REP course requirements by participating in this research. Participants were only recruited for the experiments if they were native English speakers (either monolingual or bilingual). In Experiment 1a, the data of one participant was removed from the analysis for displaying item recognition accuracy below chance level in the simultaneous

condition ($d' = -0.04$), and the data of two other participants were removed as they did not complete the experiment. In Experiment 1b, the data of one participant was removed from the analysis for displaying recognition accuracy below chance level in the blocked condition ($d' = -0.21$) and simultaneous condition ($d' = -0.37$), and the data of another participant was removed for displaying recognition accuracy below chance level in the reverse blocked condition ($d' = -0.03$).

Design. The experiments each took place over one hour-long session. Both experiments were comprised of six experimental cycles, each involving a study phase and test phase. Each study phase involved the presentation of 48 new words – 24 high frequency (HF) and 24 low frequency (LF) – on the computer screen. Half the words were presented in the bottom-left corner of the screen and the other half were presented in the top-right corner. Each word was presented in a coloured font (either green or yellow) that was redundant with the source location of the word. The assignment of the two colours to the two source locations was counterbalanced between participants.

The test phase of each study-test cycle involved making recognition and source judgements, but the order in which these were made differed across three conditions: the blocked condition, reverse blocked condition, and simultaneous condition. The blocked condition was comprised of an initial recognition test of 48 targets (24 HF, 24 LF) and 48 lures (24 HF, 24 LF), and a subsequent source memory test involving the same 48 targets. The reverse blocked condition was identical to the blocked condition, except that the source memory test preceded the recognition test. The simultaneous condition involved the presentation of 48 targets (24 HF, 24 LF) and 48 lures (24 HF, 24 LF); for each item, participants made an initial recognition judgement that was immediately followed by a source judgement for the same item (regardless of the recognition judgement made), before being presented with the next item. Throughout the first three cycles of each experiment,

participants were presented with each condition once. In the final three cycles of each experiment, participants were presented with each condition once again. The order in which each condition was presented in each half of the experiments was counterbalanced between participants.

Materials. The experimental data was collected through computer programs that were developed in PsychoPy (Peirce et al., 2019). The word list of the two experiments was initially derived from the CELEX English Database (Baayen et al., 1996). It was specified that the length of the words be limited to 5-9 letters. In order to avoid instances of highly similar words (e.g., plurals, derivational variants), words were excluded if there was another word in the overall list that began with the same three letters. From the updated word list, two new word lists were created: a list of HF words ($N = 303$) with CELEX frequencies ranging between 100-560 occurrences per million, and a list of LF words ($N = 1034$) with CELEX frequencies ranging between 1-2 occurrences per million.

Procedure. The overall length of each experiment was roughly one hour. Throughout the study phase of each experimental cycle, participants were presented with a series of individual words for 2000 ms each, separated by an interstimulus interval (ISI) of 250 ms. Each word was presented in either the bottom-left or top-right corner of the computer screen, and would appear in a coloured font that was consistent with the particular source location (either green or yellow). Upon seeing each word, participants were required to make either a “pleasantness” rating, where they pressed the “i” key if they thought the word was pleasant, or the “k” key if they thought it was unpleasant, or an “imageability” rating, where they pressed the “i” key if they thought the word was imageable, or the “k” key if they thought it was not imageable. Half of the study items from each source words would elicit the pleasantness rating, while the other half of study items from each source would elicit an imageability rating.

Following the study phase of each cycle was a card game which served as a distractor task. The card game required participants to respond to sequences of cards with rules such as, “Hit the SPACEBAR if you see two cards with the same SUIT in a row”. The card game lasted 45 seconds.

In Experiment 1a, the test phase of each study-test cycle involved participants being presented with a series of words in white font in the centre of the computer screen. When participants were required to make a recognition judgement, they were instructed to respond using a six-point confidence scale (i.e., “1”, “2”, “3”, “4”, “5”, “6”) where “1” indicated that they thought the word was definitely “OLD”, and “6” indicated that they thought the word was definitely “NEW”. When participants were required to make a source judgement, they were instructed to respond using another six-point confidence scale (i.e., “1”, “2”, “3”, “4”, “5”, “6”) where “1” indicated that they thought the word was definitely presented in the bottom-left corner, and “6” indicated that they thought the word was definitely presented in the top-right corner. The test phase of Experiment 1b was identical to Experiment 1a, except that participants were required to respond using binary yes-no judgements, such that the “1” key indicated that a word was “OLD” or from the bottom-left source, and the “0” key indicated that the word was “NEW” or from the top-right source.

In both experiments, participants were not instructed how to respond if they thought the word was a lure. Responses to test items that were faster than 200 ms resulted in a ‘TOO FAST’ warning. Responses to items that were slower than 8000 ms resulted in a ‘TOO SLOW’ warning, upon which the current trial would terminate and the next item would appear. Both types of warning messages would appear on the screen for exactly 2500 ms.

Analysis was performed on the data via a series of one-tailed Bayesian one-sample t-tests. These analyses were performed using the statistics program Just Another Stats Package (JASP; Version 0.13.1.0; Wagenmakers et al., 2018).

Results and Modelling

The following section will be presented in three key parts. First, we will present a descriptive summary of the results from each of the experiments alongside an analysis of mean source performance for unrecognized items using Bayesian statistical methods. Second, we will report a brief analysis of ROC curves for the Experiment 1a source data. Third, we will perform an analysis of source performance for unrecognized and recognized items using hierarchical Bayesian modelling, and compare the findings to those from the Bayesian statistical tests. In this section, we will describe the differing assumptions of several models that were fit to data, compare the successes of each model via a model selection metric, and then present posterior distribution plots for key parameters from the winning model. The model code we used throughout this study can be found on our OSF page (<https://osf.io/my93p/>).

Bayesian Statistical Analysis. All analyses of source accuracy for unrecognized items were conducted using Bayesian one-sample t-tests. Bayesian methods were preferred over frequentist methods due to the former's ability to provide evidence for null effects (Wagenmakers, 2007). Each test was one-tailed, such that the alternative hypothesis specified that the mean would be greater than zero. One-tailed tests were employed as there is no expectation that source memory for unrecognized items will be below chance; predictions from high-threshold and multivariate SDT models suggest that source memory for unrecognized items is either at chance or above chance respectively. The resultant Bayes factors of each test were calculated using the Jeffreys–Zellner–Siow method (JZS; Rouder et al., 2012) with normal priors ($M = 0.00$, $SD = 0.50$). Each of the reported Bayes factors (BF_{10}) represent the ratio of the alternative model likelihood to the null model likelihood, such that Bayes factors below 1.000 constitute evidence in favour of the null model, and Bayes factors above 1.000 constitute evidence in favour of the alternative model. We assume

that Bayes factors above 1, 3, 10, and 20 provide marginal, moderate, substantial, and strong evidence in favour of the alternative model respectively, while Bayes factors below 1, 1/3, 1/10, and 1/20 constitute marginal, moderate, substantial, and strong evidence in favour of the null model respectively (Jeffreys, 1961; Kass & Raftery, 1995; Lee & Wagenmakers, 2013; Stefan et al., 2019).

Table 2 displays the main findings of Experiment 1a. Means (M) and standard errors (SE) are provided for recognition hit rate (HR), recognition false alarm rate (FAR), recognition discriminability (d'), source discriminability for recognized items, and source discriminability for unrecognized items. Each of the four dependent variables are reported as a function of design type (i.e., blocked, reverse blocked, and simultaneous) and word frequency (i.e., HF words, LF words, and both word frequencies combined). We choose to report source performance using d' over percent correct, or $P(c)$, due to the fact that $P(c)$ cannot account for response bias. Accounting for bias is crucial when conditionalizing source performance on unrecognized items, as although participants are initially presented with the same number of items from each source, there is no guarantee that the number of unrecognized items will be equal from each source. While it is often preferable to report discriminability in terms of d_A instead of d' (as the latter can be affected by unequal variances; Macmillan & Creelman, 2004), it is acceptable to report our findings in terms of d' as we are primarily concerned with source discriminability and those data are typically equal-variance (Glanzer et al., 2004).

Note that in the calculations of source d' for unrecognized items, participants were excluded if they did not exhibit at least one valid response to items from each source class. As such, the sample size is reported alongside each mean and standard error. While it may be surprising to find that multiple participants were removed from each calculation of source discrimination for unrecognized items, this can be attributed to the simple fact that the mean HR was high across each condition (see Table 2). It was not uncommon for participants to exhibit the

maximum hit rate in a particular condition, and consequently, for none of the items to be unrecognized. Following this logic, it can be seen that conditions that demonstrated higher HR (e.g., LF words and the reverse blocked design) tended to decrease the effective sample size for calculating source discrimination of unrecognized items.

Table 2

Experiment 1a. Recognition and Source Performance as a Function of Design Type and Word Frequency

	Design Type								
	Blocked			Reverse Blocked			Simultaneous		
	<i>N</i>	<i>M</i>	<i>SE</i>	<i>N</i>	<i>M</i>	<i>SE</i>	<i>N</i>	<i>M</i>	<i>SE</i>
Recognition HR									
Overall	55	0.83	0.01	55	0.89	0.01	55	0.82	0.02
HF	55	0.80	0.02	55	0.86	0.02	55	0.79	0.02
LF	55	0.85	0.02	55	0.92	0.01	55	0.85	0.02
Recognition FAR									
Overall	55	0.20	0.02	55	0.17	0.02	55	0.27	0.02
HF	55	0.24	0.02	55	0.22	0.02	55	0.32	0.02
LF	55	0.15	0.02	55	0.12	0.01	55	0.22	0.02
Recognition d'									
Overall	55	1.92	0.09	55	2.39	0.10	55	1.68	0.08
HF	55	1.68	0.10	55	2.03	0.10	55	1.42	0.08
LF	55	2.22	0.10	55	2.77	0.11	55	1.98	0.11
Source d' (recognized items)									
Overall	55	0.71	0.09	55	1.12	0.10	55	0.96	0.09
HF	55	0.59	0.09	55	0.91	0.11	55	0.84	0.09
LF	55	0.82	0.11	55	1.32	0.10	55	1.09	0.11
Source d' (unrecognized items)									
Overall	53	0.13	0.08	46	0.31	0.11	52	-0.02	0.08
HF	49	0.02	0.11	40	0.21	0.13	48	0.01	0.11

LF	44	0.39	0.12	29	0.30	0.16	42	-0.03	0.10
----	----	------	------	----	------	------	----	-------	------

Note. HR, FAR, and d' estimates represent the mean of the individual participant means.

As can be seen in Table 2, recognition discriminability was also quite high across each condition. As expected, participants demonstrated a word frequency effect across each design type with LF performance consistently exceeding HF performance (Glanzer & Adams, 1985). Performance also appears to be higher in the reverse blocked condition compared to the other two design types. This is likely due to the fact that participants were exposed to each target item twice before recognition testing in this condition – once during study, and once during the source test – thus producing a strength effect. Similar trends are observed in the data for source discrimination of recognized items with participants exhibiting the usual word frequency effect and consistently higher performance in the reverse blocked condition. The latter observation is possibly due to the fact that the blocked and simultaneous conditions require participants to make a source judgement for a word after being presented with the same word detached from its original source, a practice that likely interferes with the strength of the original memory (Kim et al., 2012). This issue would not occur in a reverse blocked design as participants are presented with source judgements prior to recognition judgements.

More important are the source data for unrecognized items. In the blocked condition, while there was only marginal evidence in favour of source memory for unrecognized items ($M = 0.13$, $BF_{10} = 1.764$), the same analysis restricted to LF words yielded strong evidence ($M = 0.39$, $BF_{10} = 54.768$), and the analysis restricted to HF words revealed marginal evidence in favour of the null hypothesis ($M = 0.02$, $BF_{10} = 0.335$). In the reverse blocked condition, there was strong evidence in favour of source memory for all unrecognized items ($M = 0.31$, $BF_{10} = 20.962$) and moderate evidence for LF words ($M = 0.30$, $BF_{10} = 3.103$). Perhaps due to deficiencies in sample size, however, only marginal evidence was revealed for HF words ($M = 0.21$, $BF_{10} = 2.113$) individually. Finally, in the simultaneous condition,

The main findings of Experiment 1b are displayed in Table 3. As in Experiment 1a, recognition HRs were generally very high, particularly in the reverse blocked condition and for LF words. Consequently, there weren't many source observations for unrecognized items (see Table 3). Other trends observed in Experiment 1a were also replicated in Experiment 1b. For instance, there was an increase in item discrimination and source discrimination for recognized items in conditions involving a reverse blocked design and/or LF words.

Experiment 1b. Recognition and Source Performance as a Function of Design Type and Word Frequency

[illegible]

Overall	47	1.06	0.10	47	1.38	0.09	47	1.28	0.11
HF	47	0.84	0.12	47	1.18	0.10	47	1.03	0.12
LF	47	1.29	0.10	47	1.58	0.11	47	1.47	0.10
Source d' (unrecognized items)									
Overall	46	0.12	0.11	38	0.46	0.15	47	0.19	0.08
HF	40	0.06	0.12	31	0.36	0.15	47	0.26	0.10
LF	40	0.20	0.15	30	0.43	0.20	37	0.13	0.14

Note. HR, FAR, and d' estimates represent the mean of the individual participant means.

Findings regarding source performance for unrecognized items were generally similar to Experiment 1a, but with some peculiar differences. In the blocked condition, there was again marginal evidence in favour of the null hypothesis for HF items ($M = 0.06$, $BF_{10} = 0.455$) and marginal evidence for all items ($M = 0.12$, $BF_{10} = 0.828$). This time, however, only marginal evidence was established for LF items as well ($M = 0.20$, $BF_{10} = 1.337$). In the reverse blocked condition, there was again strong evidence in favour of source memory for unrecognized items ($M = 0.46$, $BF_{10} = 26.849$) and **moderate** evidence for LF items specifically ($M = 0.43$, $BF_{10} = 4.218$). One difference, however, was that **moderate** evidence was established for HF items ($M = 0.36$, $BF_{10} = 6.330$). Interestingly, the findings in the simultaneous condition were largely the opposite of Experiment 1a. Instead of yielding **moderate** evidence in favour of the null hypothesis, the Bayesian one-sample t-tests revealed substantial evidence in favour of source memory for all unrecognized items ($M = 0.19$, $BF_{10} = 5.649$) and for HF items specifically ($M = 0.26$, $BF_{10} = 13.184$), as well as marginal evidence for LF items ($M = 0.13$, $BF_{10} = 0.740$).

Another intriguing aspect of the data was that the percent of correct source responses for unrecognized items was consistently higher for the top-right source than the bottom-left source. In Experiment 1a, the source P(c) for unrecognized items in the “bottom-left” and “top-right” sources respectively were .47 and .59 in the blocked design, .53 and .59 in the

reversed design, and .26 and .73 in the simultaneous design. In Experiment 1b, these values were .51 and .54 in the blocked design, .53 and .64 in the reversed design, and .29 and .77 in the simultaneous design. While this appears to suggest that top-right items were better encoded than bottom-left items (perhaps due to an attentional bias), this is unlikely the case given item recognition d' was the same for each source across all conditions. Specifically, the mean item recognition d' for bottom-left and top-right sources respectively was 1.11 and 1.12 in Experiment 1a, and 1.08 and 1.05 in Experiment 1b. If there was an attentional bias to the top-right source, then you would also see higher values of recognition discriminability for those items. Instead, it is possible that the apparent response bias to the top-right source resulted from participants favouring that response when providing guesses. This would make sense considering that participants used the same key(s) for “NEW” recognition responses and “top-right” source responses.

Put together, what do the experiments reveal regarding the dependency of source discrimination in the absence of recognition on experimental design? For the blocked condition, there was clear evidence of source memory for unrecognized items for LF words, but not so for HF words. For the reverse blocked condition, however, clear and strong evidence was established in favour of the effect across the board. Interestingly, the results from the simultaneous condition were inconsistent, with there being no evidence for the effect in Experiment 1a, but clear evidence for the effect in Experiment 1b. Overall, it is fair to say that the results verify the notion that source memory for unrecognized items is largely dependent on experimental design, as there was stronger and more consistent evidence of the effect in the two blocked designs than the simultaneous design.

Analysis of Receiver Operating Characteristic Curves. Along with the Bayesian statistical analysis, it is useful to perform an analysis of the ROC data from Experiment 1a. Note that ROC data cannot be calculated for Experiment 1b as participants did not provide

confidence judgements. The group averaged ROC data for source memory in Experiment 1a are displayed in Figure 4. Note that each ROC curve represents a particular old-new rating that participants provided to items during the item recognition phase. Across the blocked, reverse blocked, and simultaneous conditions, one can see that the source ROC curves for targets rated “1” (definitely “OLD”) strongly deviate from the diagonal, indicating that source discriminability for highly recognized items is well above chance. However, while the source ROCs for items recognized with lower confidence (“2” and “3”) are strongly curved in the reverse blocked condition, they adhere closely to the diagonal in the blocked and simultaneous conditions, indicating chance-level source discriminability for those items. For targets rated “NEW” with low confidence (“4” and “5”), the source ROCs are strongly curved in the reverse blocked condition, slightly curved in the blocked condition, and linear in the simultaneous condition.

Remarkably, the source ROCs for targets rated “6” (definitely “NEW”) appear to be strongly curved across all three conditions. In fact, only the targets rated “1” have source ROCs that deviate further from the diagonal in each condition. This suggests that source memory is above-chance for unrecognized items across all three conditions (although it is stronger in the two blocked condition). While it may seem surprising that source memory is stronger for items that are unrecognized with higher confidence **than** with lower confidence, the same phenomena can be observed in the ROC data of Osth et al. (2018, Exp. 3) and Starns et al. (2013, Exps. 2a & 2b). Even for investigations that did not report above-chance source memory for unrecognized items, it can be seen that the source ROCs for targets unrecognized with high confidence are more curved than those for targets unrecognized with low confidence (Onyper et al., 2010, Exp. 2; Yonelinas, 1999, Exp. 2). Furthermore, it is unlikely that the phenomenon occurred due to response switching – with participants accidentally responding “definitely NEW” instead of “definitely OLD” – as recognition

performance was highest at the two ends of the confidence scale (i.e., participants did not appear to make many errors when responding confidently). Specifically, the proportion of correct responses when responding maybe, probably, and definitely “NEW” was 0.76, 0.82, and 0.90 respectively. Note as well that the definitely “OLD” and “NEW” responses were maximally far apart on the keyboard, so it is unlikely that motor errors were responsible. That said, it is also important to acknowledge the impact that noise may have had on these ROC curves. Given that “definitely NEW” judgements were rarely prescribed to targets in our data, many participants exhibited too few of these items to accurately calculate source discriminability. For instance, in the reverse blocked condition, source discriminability for “definitely NEW” targets could be only be calculated for 19 participants, and the median number of relevant trials that these participants exhibited was only 4.

Figure 4

Group Averaged ROC Data for Experiment 1a for Source Memory

Note. Group averaged ROC data for Experiment 1a for source memory in the blocked (left panel), reverse blocked (middle panel), and simultaneous conditions (right panel). ROC curves are presented for each old-new rating that participants provided to items in the recognition phase (“1” = definitely “OLD”, “6” = definitely “NEW”).

Hierarchical Bayesian Analysis. As mentioned earlier, one of the primary goals throughout the present investigation is to investigate empirical findings using hierarchical Bayesian SDT models. This goal is of particular importance in Experiments 1a and 1b given the generally low number of unrecognized items per participant and consequently heightened degree of uncertainty in the participant-level estimates of source memory for unrecognized items. As can be seen in Figure 5 that shows scatterplots of item recognition hit rate against source d' for unrecognized items, the degree of uncertainty in the source d' estimates gets increasingly high as HR increases. While traditional forms of statistical analysis (e.g., the Bayesian t-test of Rouder et al., 2012) use point estimates for the means of individual participants, Bayesian models account for uncertainty at the participant-level and quantify it as posterior distributions. Furthermore, hierarchical Bayesian models (as opposed to standard Bayesian models) incorporate shrinkage, meaning that highly uncertain participant-level parameters (particularly outliers) can be adjusted so that they are closer to the group-level parameters (Boehm et al., 2018). Given that the Bayesian t-test analyses we reported above relied on point estimates for each participant, it is possible that the conclusions we gleaned will differ when the Bayesian SDT model is applied. Recall that Osth et al. (2018) measured above-chance source accuracy for unrecognized items with a hierarchical Bayesian SDT model, when only chance-level accuracy was observed in their analysis of point estimates.

Figure 5

Relationship Between Participant Measures of Item Recognition HR and Source Discriminability for Unrecognized Items

Note. Relationship between participant measures of item recognition HR and source d' for unrecognized items in Exp. 1a (left panel) and Exp. 1b (right panel). Data from the blocked, reverse blocked, and simultaneous conditions are represented by blue, red, and green dots respectively.

Model Selection. Given that cognitive models inherently make assumptions about the nature of the processes they are trying to capture, it is good practice to devise multiple different models and compare their ability to fit the data before drawing strong conclusions from the fit of one particular model. For Experiments 1a and 1b, five different SDT models were devised to investigate source data for both unrecognized items and recognized items. All models included a criterion (c) for each of the available response options (five criteria in Exp. 1a, one criterion in Exp. 1b) and allowed the criteria to vary across the three design conditions, thus accounting for 15 free parameters in the Experiment 1a models and 3 free parameters in the Experiment 1b models. The models differed with regards to their assumptions about the discrimination (d') parameter. The first model fit to each experiment fixed the d' parameter to zero across all conditions (the $d' = 0$ model), and as such, only required the estimation of decision criteria. The second model included a single free

parameter to represent d' across all conditions (the Single d' model). The third and fourth models further allowed d' to vary across the three design conditions ($d' \sim D$, where we use $\sim$ to mean “varies as a function of”) and the two word-frequency classes ($d' \sim WF$) respectively, while the fifth model allowed d' to vary across both ($d' \sim D, WF$).

Each of the five models was fit within a hierarchical Bayesian framework with sampling governed by the Stan programming language (Gelman et al., 2015). The use of Stan was preferred to other sampling methods, such as JAGS (Plummer, 2003), as it is more efficient at dealing with correlated parameters, which can occur frequently in cognitive models (Annis et al., 2017). For each model, the number of chains was set to equal three times the number of parameters. Sampling began after a burn-in period of 6,000 iterations and involved one in every 20 samples being collected until there was a total of 2,000 samples in each chain. The prior distributions used for each of the model parameters were relatively non-informative. Further details of the priors we employed are included in Appendix A, details regarding model convergence are provided in Appendix B, and posterior checks for the preferred models are provided in Appendix C.

To compare the fits of each model, we used the widely applicable information criterion (WAIC; Watanabe, 2010). WAIC is a model selection technique that quantifies the strength of a model in fitting a dataset, while subtracting a quantity that corresponds to the complexity of the model. Smaller values of WAIC indicate that a model is more parsimonious, meaning that it provides a good fit of the data while avoiding unnecessary complexity. Since WAIC is measured on a log likelihood scale, differences in WAIC that are greater than 10 are conventionally considered large. For ease of interpretation, we will not report the raw WAIC values for each model, but instead report the change in WAIC for each model ($\Delta WAIC$) relative to the winning model. Given that WAIC values can often be very close for competing models, we will report weights (Wagenmakers & Farrell, 2004)

alongside the ΔWAIC values that represent the conditional probabilities of each model. Table 4 reports the ΔWAIC values and WAIC weights $w(\text{WAIC})$ for each model across both experiments, and for both unrecognized and recognized items.

Table 4

ΔWAIC and $w(\text{WAIC})$ Values for Each Hierarchical Bayesian SDT Model Applied to the Unrecognized and Recognized Items from Experiments 1a and 1b

Model	Experiment 1a					Experiment 1b				
	N	Unrecognized		Recognized		N	Unrecognized		Recognized	
		Δ WAIC	w (WAIC)	Δ WAIC	w (WAIC)		Δ WAIC	w (WAIC)	Δ WAIC	w (WAIC)
$d' = 0$	15	24.53	0.00	3567.66	0.00	3	36.46	0.00	3140.65	0.00
Single d'	16	1.77	0.19	182.56	0.00	4	13.63	0.00	51.06	0.00
$d' \sim \mathbf{D}$	18	9.91	0.00	122.51	0.00	6	0.00	0.87	66.57	0.00
$d' \sim \mathbf{WF}$	17	0.45	0.36	48.39	0.00	5	14.93	0.00	0.00	1.00
$d' \sim \mathbf{D, WF}$	21	0.00	0.45	0.00	1.00	9	3.88	0.13	46.06	0.00

Note. Winning models are highlighted in bold. N = number of parameters per participant.

The results in Table 4 indicate a strong rejection of the $d' = 0$ model for both Experiments 1a and 1b, indicating that there is source memory for unrecognized items. For Experiment 1a, the most preferred model in capturing unrecognized and recognized items was the model that allowed d' to vary across design conditions and frequency classes. For Experiment 1b, the preferred models for capturing unrecognized and recognized items were the models that allowed d' to vary solely across design and solely across word frequency respectively. Below we present plots of source data for unrecognized and recognized items alongside posterior distributions of the group means of the d' parameters from the preferred models, each represented by 95% highest density intervals (HDIs). Note that while important insights are derived from the following analyses of posterior distributions, our conclusions from the hierarchical Bayesian modelling procedure are primarily derived from the model

selection procedure (Table 4). This is because our analysis of posterior distributions is only at the group level, while the model selection procedure reflects model fit at both the group level and participant level.

Posterior Distributions of Source Discriminability. Figure 6 presents source discrimination for unrecognized and recognized items in Experiment 1a as a function of design type (i.e., blocked, reverse blocked, simultaneous) and word frequency (i.e., HF, LF). While the Bayesian model measurements of source d' for recognized items closely mirror the data across all conditions, the measurements of source d' for unrecognized items show some deviation from the data. Most importantly, it can be seen that the 95% HDIs of the group mean d' parameters are consistently higher than the means estimated directly from the data. For instance, while the data suggests that accuracy is at chance for HF items in the simultaneous condition, the model suggests that accuracy is comfortably above chance (zero is not included within the 95% HDIs). Furthermore, while the Bayesian one-sample t-tests were unable to establish substantial evidence that accuracy in the reverse blocked condition was above chance for HF items, the 95% HDIs of the model estimates suggest that accuracy is above chance in the condition.

Figure 6

Mean Source Discriminability for Unrecognized and Recognized Items in Experiment 1a

Note. Mean source discriminability (d') for unrecognized and recognized items in Experiment 1a as a function of design type (blocked, reverse blocked, simultaneous) and word frequency (HF, LF).

Posterior distributions of the group means of the d' parameters from the winning model ($d' \sim D$, WF) are represented by 95% HDIs.

Figure 7 presents source discrimination for unrecognized and recognized items in Experiment 1b. As with the previous experiment, the Bayesian model measurements are highly consistent with the data for recognized items, but show some deviation for unrecognized items. With the exception of the blocked condition, the means of the model's group-level d' distributions are all higher than the mean based on participants' point estimates.

Figure 7

Mean Source Discriminability for Unrecognized and Recognized Items in Experiment 1b

Note. Mean source discriminability (d') for unrecognized and recognized items in Experiment 1b.

Posterior distributions of the group means of the d' parameters from the winning models ($d' \sim D$ for unrecognized items, $d' \sim WF$ for recognized items) are represented by 95% HDIs.

For the two blocked conditions, the results from the Bayesian modelling are fairly consistent with the results from the Bayesian statistical tests, suggesting that source memory is generally above chance in the absence of recognition. As alluded to earlier, the higher estimates for the reversed condition likely results from the fact it requires participants to make source judgements prior to recognition judgements. Since recognition tests involve the presentation of target items that are detached from their original source manipulations, observing them prior to source retrieval would generate additional interference and weaken source memory (Kim et al., 2012).

Interestingly, the Bayesian model estimates also suggest that the effect is considerably higher in the simultaneous condition than what was evident from the Bayesian statistical tests. This observation is consistent with Osth et al. (2018) who observed chance-level accuracy with their analysis of point estimates, but measured above-chance accuracy with a

hierarchical Bayesian SDT model. It is probable that the statistical tests we employed were mostly insensitive to the effect in the simultaneous condition due to multiple participants exhibiting too few observations of unrecognized items to reliably estimate d' , and therefore exhibiting estimates that were near zero. The hierarchical Bayesian model is able to correct this issue via shrinkage, whereby uncertain (and likely near-zero) participant estimates are adjusted up toward the group mean.

The observation of above-chance source accuracy for unrecognized items in the simultaneous condition challenges the notion that the effect is dependent on a particular experiment design and instead suggests that it is a legitimate psychological phenomenon. That said, the majority of the evidence for the effect stems from the two blocked conditions (particularly the reverse blocked condition). One explanation is that participants are more motivated to perform source retrieval for items they don't recognize in these conditions, given that the recognition trials do not directly precede the source trials. On the other hand, one could argue that evidence for the effect in the blocked conditions simply results from changes in memory state between blocks. In other words, apparent source memory for unrecognized items in the blocked conditions may result from a subset of unrecognized items during recognition being implicitly recognized during source retrieval.

Experiments 2a and 2b

In light of the concern that source accuracy for unrecognized items in the blocked and reversed conditions resulted from changes in memory between testing blocks, two new experiments were devised to verify whether this was indeed the case. The two experiments (one using confidence judgements and the other using yes-no judgements) substituted the source memory judgements in the aforementioned blocked and reverse blocked designs with a *remember-forget* task. This task, inspired by Koop et al. (2014), is distinct from a standard recognition task in that participants are only presented with targets (not targets and lures), and

are simply asked whether they remember or forget seeing each item on the study list. In doing so, memory for study list items can be directly compared between different blocks. If the vast majority of unrecognized items are labelled as “forgotten” in the remember-forget task, then that would suggest that item memory is consistent throughout testing. On the other hand, if a sizeable proportion of unrecognized items are labelled as “remembered”, then that would suggest that item memory is prone to fluctuations from block to block, and that source performance for unrecognized items could actually reflect source performance for items that were at one point unrecognized, but were recognized at the time of source retrieval.

While it may appear that the same analysis could be achieved by utilising two distinct blocks of recognition testing, it is important to note that only target items are presented during the source phase of the blocked conditions, and as such, performance in a dual-recognition design may reflect the presence of several new lures, and not be directly comparable to the standard designs. Essentially, the remember-forget design is the most analogous to the standard source memory design, the only difference being the task itself. Furthermore, while it may seem strange to employ a non-recognition task in the hopes of extracting recognition-like information, it has been shown that presenting participants with exclusively targets or lures, and not both, can produce data that is remarkably similar to recognition (Koop et al., 2014; Wallace, 1982).

Method

Participants. Both experiments were approved by the Human Research Ethics Committee of The University of Melbourne (Ethics ID no. 12033). For Experiment 2a, the participants were $N = 33$ English-speaking students aged 18 or older at the University of Melbourne, none of which had partaken in the first set of experiments. For Experiment 2b, the participants were $N = 34$ students from the same population, none of which had partaken in any of the other experiments. As with the first set of experiments, all participants across

Experiments 2a and 2b were recruited through the Melbourne School of Psychological Sciences REP and received credit towards fulfilling their course requirements by participating. In Experiment 2a, the data of two participants were removed from the analysis for displaying item recognition discriminability near chance ($d' < .35$) in the blocked and reversed conditions respectively, and the data of three other participants were removed for displaying near-chance recognition discriminability in both conditions. In Experiment 2b, the data of one participant was removed from the analysis for displaying near-chance recognition discriminability in the blocked condition, and the data of three other participants were removed for displaying near-chance recognition discriminability in both conditions.

Design. The design aspects of Experiments 2a and 2b were largely carried over from the first set of experiments. The key difference was that during test, participants were presented with recognition and “remember-forget” judgements, as opposed to recognition and source judgements. The order of these judgements in the test phase was dictated by two separate conditions: blocked and reverse blocked. In the blocked condition, participants were presented with an initial recognition test of 48 targets (24 HF, 24 LF) and 48 lures (24 HF, 24 LF), and a subsequent remember-forget task involving the same 48 targets. In the reverse blocked condition, participants completed the remember-forget task prior to the recognition task. Across the six study-test cycles, participants completed each condition three times in alternating order. In each experiment, the condition of the first cycle (either blocked or reverse blocked) and the assigned colours of the two source locations (either green-yellow or yellow-green) were counterbalanced between participants.

Materials. The experimental data was collected through online programs that were created using the JavaScript programming language with the jsPsych library. The word list of the two experiments was identical to that of Experiments 1a and 1b.

Procedure. As with the design, the majority of the procedural details of Experiments 2a

and 2b were carried over from Experiments 1a and 1b. The sole difference was the presence of the remember-forget task, as opposed to the source task. In Experiment 2a, when participants were required to make a remember-forget judgement, they were instructed to respond using a six-point confidence scale (i.e., “1”, “2”, “3”, “4”, “5”, “6”) where “1” indicated that they strongly remembered seeing the target item on the study list, and “6” indicated that they had strongly forgotten seeing the target item on the study list. In Experiment 2b, participants provided responses using binary yes-no judgements, such that the “1” key indicated that a word was remembered, and the “0” key indicated that the word was forgotten.

In both experiments, participants were clearly instructed that targets and lures would appear in the recognition task but that only targets would appear in the remember-forget task. For the remember-forget task, they were instructed to respond with the utmost honesty, such that a forgotten target would be labelled as “forgotten” and not mislabelled as “remembered” as a way to cover up imperfect memory.

Results and Simulations

Figure 8 presents the proportions of targets in Experiments 2a (top row) and 2b (bottom row) that were recognized or unrecognized in the recognition phase and remembered or forgotten in the remember-forget phase. The left panels display results from the blocked condition and the right panels display results from the reversed condition. Across both experiments and both conditions, the most common outcome for a target was to be both recognized in the recognition phase and remembered in the remember-forget phase. However, in all conditions, a sizable proportion of recognized targets were contradictorily labelled as “forgotten” in the alternate blocks of trials. Even more strikingly, of the items that were unrecognized in the recognition phase, there was a consistently equal chance of those items being labelled as “remembered” or “forgotten” in the alternate block of trials. If

participants' memory of items was consistent across blocks of testing, then one would expect the proportions of recognized/forgotten targets and (especially) unrecognized/remembered targets to be near zero. While cases of recognized/forgotten and remembered/unrecognized targets in the blocked and reversed conditions respectively could be attributed to interference-based forgetting from one block to the next, the cases of unrecognized/remembered and forgotten/recognized targets in the blocked and reversed conditions respectively could be attributed to participants updating their memory traces for target items after witnessing them for a second time in the first block of testing. For instance, participants may claim to remember seeing a previously unrecognized target on the study list if they have a strong memory of seeing it during the recognition phase, but confuse the context of that memory with the study-list context.

Figure 8

Results of Experiments 2a and 2b

Note. Proportions of target items that are recognized/unrecognized and remembered/forgotten as a

function of design type (blocked or reverse blocked) and response modality (confidence or yes-no ratings) in Experiments 2a and 2b. Error bars represent the standard error of the mean.

To determine whether these occurrences of contradictory responses between the recognition and remember-forget phases would be sufficient to produce the degree of source memory for unrecognized items that we observed in the two blocked conditions of Experiment 1a and 1b, we performed the following simulations. For each participant in Experiments 1a and 1b, we generated expected values of source accuracy for recognized and unrecognized targets in the blocked and reverse blocked conditions that were based upon the mean proportions of items that flipped responses between blocks in Experiments 2a and 2b. For instance, if there was a 50/50 chance that an unrecognized target in the blocked condition of Experiment 2a would subsequently be labelled as “remembered” or “forgotten”, then each unrecognized item in the blocked condition of Experiment 1a would be reclassified as either recognized or unrecognized with 50% probability during the source phase. For each item reclassified as recognized, an artificial source judgement would be generated with a probability of being correct that matched the participant's source accuracy for recognized items in the simultaneous condition. Likewise, for items that remained classified as unrecognized, an artificial source judgement would be generated with a probability of being correct that matched the participant's source accuracy for unrecognized items in the simultaneous condition. Calculations of the same kind were performed for recognized items (with each item being reclassified as either recognized or unrecognized), for the reverse blocked condition (with source accuracy in the first block being simulated), and for Experiment 1b (with rates of response flipping dictated by the Experiment 2b data).

Figure 9 presents the empirical and simulated source accuracy for unrecognized and recognized items across blocked and reverse blocked conditions in Experiments 1a and 1b.

For the source accuracy for unrecognized items in each experiment and condition, it can be seen that the simulated values reflecting the response flipping from Experiments 2a and 2b are not only sufficient to match the empirical values, but in fact, consistently exceed them. This suggests that source memory for unrecognized items in the blocked and reverse blocked conditions may have primarily resulted from previously unrecognized items becoming recognized, as opposed to reflecting source accuracy for items that were legitimately unrecognized. At the very least, the results of Experiments 2a and 2b suggest that the two blocked designs of Experiments 1a and 1b are not suitable for investigating source memory for unrecognized items, as the assumption that memory would be consistent from one block to another did not hold under scrutiny. Instead, it appears that blocked designs induce considerable variability in aspects such as learning and criterion position. Regarding the greater literature, the failed validation of the blocked designs is especially concerning given that source memory for unrecognized items is typically reported in these designs.

Figure 9

Simulations of Performance in Experiments 1a and 1b Based on Results from Experiments 2a and 2b

Note. Source memory accuracy (percent correct) for recognized and unrecognized items as a function of design type (blocked or reverse blocked) and response modality (confidence or yes-no ratings) in Experiments 1a and 1b. Blue bars represent empirical estimates of source accuracy and red bars represent simulated estimates of source accuracy based upon the rates of response flipping between blocks in Experiments 2a and 2b. Error bars represent the standard error of the mean.

Experiments 3a and 3b

In light of the findings from Experiments 2a and 2b that blocked designs are not suitable for investigating source memory for unrecognized items, two new experiments are devised (one with confidence judgements and the other with yes-no judgements) that focus solely on the simultaneous condition. While the Bayesian modelling exercise for the first set of experiments yielded convincing evidence for the effect in the simultaneous condition, the conclusions from the Bayesian statistical tests were mixed. As such, there remains a concern that the findings from those experiments were overly dictated by deficiencies in the

participant-level data (i.e., too few unrecognized items per participant) and group-level data (i.e., too small a sample size). To alleviate these issues, Experiments 3a and 3b present participants with three times the number of targets in the simultaneous condition than in Experiments 1a and 1b (288 vs. 96 respectively). Furthermore, we intend to collect roughly three times the number of participants for each experiment than in Experiments 1a and 1b. By collecting rich amounts of data per participant in the simultaneous condition, and incorporating a considerably larger sample than in previous experiments, Experiments 3a and 3b aim to establish sufficiently reliable and valid measures of source memory for unrecognized items.

Method

Participants. Both experiments were approved by the Human Research Ethics Committee of The University of Melbourne (Ethics ID no. 12033). In Experiments 3a and 3b, data was collected from three separate pools of participants: (a) English-speaking University of Melbourne students aged 18 or older that were recruited through the REP; (b) English-speaking individuals aged 18 or older that were recruited through the Melbourne School of Psychological Sciences Reimbursed Participant System (RPS); and (c) English-speaking individuals from the USA aged 18 or older that were recruited through Amazon's Mechanical Turk (mTurk) service. All REP participants received 1-hour of credit towards their REP course requirement and all RPS participants were reimbursed AU\$10 for their efforts. The mTurk workers were only recruited for the study if they exhibited a human intelligence task (HIT) approval rate above 95% and a number of HITs approved above 100. Upon completion of the task, each mTurk participant was reimbursed US\$10. For Experiment 3a, the participants were $N = 188$ individuals (140 from REP, 20 from RPS, and 28 from mTurk), none of which who had participated in the first two sets of experiments. For Experiment 3b, the participants were $N = 186$ individuals (141 from REP, 22 from RPS, and 23 from mTurk),

none of which who participated in any of the other experiments. In Experiment 3a, the data of 35 participants was removed (24 from REP, 2 from RPS, 9 from mTurk) for exhibiting item recognition discriminability that was near chance ($d' < .35$). In Experiment 3b, the data of another 30 participants was removed (17 from REP, 3 from RPS, 10 from mTurk) for exhibiting near-chance recognition discriminability.

Although the proportion of discarded participants across both experiments was considerably higher (17%) than across the first set of experiments (5%), this was most likely due to the nature of testing and the participant pool, as opposed to experiment-related issues (especially given that the experiments were virtually identical). Note that all data in the first set of experiments were collected within quiet laboratories at the University of Melbourne, with participants being supervised at all times. In contrast, all participants from Experiments 3a and 3b completed their online tasks in uncontrolled and unsupervised environments, which may have enabled some to rush through the tasks with no/minimal effort. It is important to consider that the mean proportion of rushed recognition responses ($RT < 500$ ms) amongst the culled participants from Experiments 3a and 3b was considerably higher than that of the accepted participants (.29 and .02 respectively).

Design. The design aspects of Experiments 3a and 3b were almost identical to the first set of experiments. While participants in Experiment 1a and 1b completed two cycles each of the blocked, reverse blocked and simultaneous conditions, participants in Experiments 3a and 3b completed six cycles of the simultaneous condition.

Materials. As with Experiments 2a and 2b, the experimental data was collected through online programs that were created using the JavaScript programming language with the jsPsych library. The same word lists were used as in the first two sets of experiments.

Procedure. The procedural aspects of Experiments 3a and 3b were also identical to the first set of experiments. While Experiment 3a utilised six-point confidence scales for

recognition and source judgements, Experiment 3b utilised binary yes-no judgements.

Results and Modelling

As with the results of Experiments 1a and 1b, this following section will be divided into key parts. First, we will present descriptive statistics for Experiments 3a and 3b, and then a Bayesian statistical analysis of source memory for unrecognized items. Second, an ROC analysis will be reported for the Experiment 3a data. Lastly, we will present a analysis of source memory for recognized and unrecognized items using hierarchical Bayesian methods, including a model selection exercise and presentation of posterior distributions for the winning model's d' parameters.

Bayesian Statistical Analysis. Table 5 displays descriptive statistics for recognition HR, FAR, recognition d' , and source d' for recognized and unrecognized items in Experiments 3a and 3b. As before, participants were excluded from the group means of unrecognized source d' if they exhibited less than one response to items from either source class. Given that participants were presented with three times the number of trials in Experiments 3a and 3b than in the simultaneous condition of Experiments 1a and 1b, far fewer needed to be excluded (see Tables 2 & 3). Furthermore, the greater number of trials per participant and the larger samples gave way to more robust estimates for each of the dependent variables (as reflected by the lower values of SE). As expected, participants also exhibited a strong word frequency effect for recognition d' and source d' for recognized items, with LF words outperforming HF words (Glanzer & Adams, 1985).

Table 5

Experiments 3a and 3b. Recognition and Source Performance as a Function of Word Frequency

Simultaneous Condition

	Exp. 3a (Confidence Ratings)			Exp. 3b (Yes-No Ratings)		
	<i>N</i>	<i>M</i>	<i>SE</i>	<i>N</i>	<i>M</i>	<i>SE</i>
Recognition HR						
Overall	153	0.78	0.01	156	0.75	0.01
HF	153	0.77	0.01	156	0.73	0.01
LF	153	0.79	0.01	156	0.76	0.01
Recognition FAR						
Overall	153	0.25	0.01	156	0.17	0.01
HF	153	0.30	0.01	156	0.21	0.01
LF	153	0.20	0.01	156	0.14	0.01
Recognition d'						
Overall	153	1.61	0.05	156	1.77	0.06
HF	153	1.43	0.05	156	1.61	0.06
LF	153	1.84	0.07	156	1.99	0.07
Source d' (recognized items)						
Overall	153	1.00	0.05	156	1.32	0.06
HF	153	0.89	0.05	156	1.21	0.06
LF	153	1.11	0.06	156	1.44	0.06
Source d' (unrecognized items)						
Overall	152	0.09	0.04	156	0.06	0.03
HF	152	0.10	0.05	156	0.06	0.04
LF	151	0.09	0.04	155	0.11	0.04

Note. HR, FAR, and d' estimates represent the mean of the individual participant means.

As with Experiments 1a and 1b, all analyses of source accuracy for unrecognized items were conducted using one-tailed Bayesian one-sample t-tests with normal priors ($M = 0.00$, $SD = 0.50$). In Experiment 3a, moderate evidence was established in favour of source memory for unrecognized items for LF items ($M = 0.09$, $BF_{10} = 3.073$), similar evidence was established for HF items ($M = 0.10$, $BF_{10} = 2.772$), and moderate evidence was established

for both frequencies combined ($M = 0.09$, $BF_{10} = 6.841$). In Experiment 3b, moderate evidence in favour of source memory for unrecognized items was also established for LF items ($M = 0.11$, $BF_{10} = 7.735$), but the evidence for HF items ($M = 0.06$, $BF_{10} = 0.697$) and both frequencies combined ($M = 0.06$, $BF_{10} = 1.220$) was inconclusive. Put together, Experiments 3a and 3b provide moderate evidence that there is source memory for unrecognized LF words, but inconclusive evidence that there is an effect for HF words. This contradicts Experiment 1a that found moderate evidence against any effect in the simultaneous condition, as well as Experiment 1b that found evidence for the effect, but primarily for HF words.

Analysis of Receiver Operating Characteristic Curves. As with the confidence ratings of Experiment 1a, we have also performed an ROC analysis of the confidence ratings of Experiment 3a. Figure 10 presents the group averaged ROC data for source memory in Experiment 3a, with each curve representing a different old-new rating that participants provided to items during recognition. As with the simultaneous condition in Experiment 1a (see Figure 4), the source ROC curve for targets rated “1” (definitely “OLD”) strongly deviates from the diagonal, indicating that source discriminability for strongly recognized items is well above chance. Unlike Experiment 1a though, the source ROCs for items recognized with lower confidence (“2” and “3”) are also strongly curved in Experiment 3a, indicating above-chance source discriminability for those items. Interestingly, while the source ROCs in Experiment 1a for targets rated “NEW” with low confidence (“4” and “5”) and high confidence (“6”) were linear and curved respectively (indicating chance-level and above-chance accuracy respectively), the source ROCs in Experiment 3a for targets rated “NEW” were curved for both the most confident ratings (“5” and “6”), and only linear for the least confident rating (“4”).

As in Experiment 1a, the finding of source memory for the most confidently

unrecognized items is peculiar, but is also consistent with the ROC data of prior investigations (Osth et al., 2018; Starns et al., 2013). Participants also exhibited the highest levels of recognition accuracy for the most confident response options, indicating that an abundance of accidental response switching for those ratings was unlikely. Furthermore, while there was a concern that the Experiment 1a source ROCs for “definitely NEW” items may have been impacted by sample noise, this is unlikely to be the case for the Experiment 3a ROC data, given that the “6” and “5” curves were calculated from a larger number of participants (131 and 134) and a higher median number of items per participant (20 and 23.5) than the “4” curve (128 participants, 15 trials median).

Figure 10

Group Averaged ROC Data for Experiment 3a for Source Memory

Note. ROC curves are presented for each old-new rating that participants provided to items in the recognition phase (“1” = definitely “OLD”, “6” = definitely “NEW”).

Hierarchical Bayesian Analysis. As with the first set of experiments, the data of Experiments 3a and 3b will also be subjected to a hierarchical Bayesian analysis. Although

this analysis was more crucial for the first set of experiments – given that there were roughly three times fewer participants and three times fewer trials in the simultaneous condition – it will still be useful for the current experiments as the proportions of targets that were unrecognized were still low (22% in Exp. 3a and 25% in Exp. 3b).

Model Selection. For the two new experiments, three different SDT models were used to investigate source performance for recognized and unrecognized items. Each of these models were previously used to investigate the Experiment 1a and 1b data: the $d' = 0$ model, the Single d' model, and the $d' \sim \text{WF}$ model. As before, each of the models were fit within a hierarchical Bayesian framework with sampling in Stan (details regarding sampling and priors were unchanged). The model convergence details are provided in Appendix B and posterior checks for the preferred model are provided in Appendix C. As before, each model was compared using WAIC (Watanabe, 2010). Table 6 presents the ΔWAIC and $w(\text{WAIC})$ values of each model across Experiments 3a and 3b for unrecognized and recognized items.

Table 6

ΔWAIC and $w(\text{WAIC})$ Values for Each Hierarchical Bayesian SDT Model Applied to the Unrecognized and Recognized Items from Experiments 3a and 3b

Experiment 3a						Experiment 3b				
		Unrecognized		Recognized		Unrecognized		Recognized		
Model	N	Δ WAIC	$w(\text{WAIC})$	Δ WAIC	$w(\text{WAIC})$	N	Δ WAIC	$w(\text{WAIC})$	Δ WAIC	$w(\text{WAIC})$
$d' = 0$	5	24.84	0.00	11400.51	0.00	1	10.07	0.01	12755.32	0.00
Single d'	6	9.48	0.01	385.52	0.00	2	4.09	0.11	94.86	0.00
$d' \sim \text{WF}$	7	0.00	0.99	0.00	1.00	3	0.00	0.88	0.00	1.00

Note. Winning model is highlighted in bold. N = number of parameters per participant.

For both Experiments 3a and 3b, the ΔWAIC values in Table 6 indicate a strong rejection of the $d' = 0$ model, suggesting that there is source memory for unrecognized items.

For both experiments, the most favoured model in capturing source performance for unrecognized and recognized items was the model that allowed d' to vary between frequency classes. Throughout the remainder of this section, we will present plots displaying the posterior distributions of the group means of the d' parameters from the preferred model (represented by 95% HDIs) alongside the source data for unrecognized and recognized items.

Posterior Distributions of Source Discriminability. Figure 11 presents source d' for unrecognized and recognized items in Experiments 3a and 3b as a function of word frequency. For both recognized and unrecognized items in Experiment 3a, there is a relatively high degree of correspondence between the Bayesian model measurements of source d' and the data. This degree of correspondence, which was not observed in Experiment 1a, is likely due to there being much more data per participant in Experiment 3a. The most notable aspect of Figure 11, however, is that the 95% HDIs of the group mean d' parameters for unrecognized items do not include zero (although it is very close for HF words), suggesting that source memory for unrecognized items is slightly above chance. This reinforces both the findings from the Bayesian statistical analysis and the WAIC analysis.

Regarding Experiment 3b, while there is high correspondence between the Bayesian model measurements and data for recognized items, there is some slight deviation for unrecognized items. Namely, the means of the model's group-level d' distributions are both lower than the point estimates from the data. This is surprising given that the model estimates were typically higher than the point estimates from the data in the analyses of Experiments 1a and 1b. Most notably, the 95% HDIs for both of the model's group-level d' distributions include zero. This appears to contradict the WAIC analysis that revealed a strong rejection of the $d' = 0$ model. It is important to note, however, that while the 95% HDIs in Figure 11 represent the group mean of the d' parameters, the WAIC analysis reflects model fit at both the group level and participant level. With regards to the participant estimates of d' , the $d' = 0$

model assumes that all participants exhibit a value of d' that is exactly zero; in contrast, the Single d' model and $d' \sim \text{WF}$ models allow d' to vary across participants. It is likely that the WAIC analysis favoured the latter models due to a proportion of participants benefiting from having $d' > 0$. So, while we cannot be confident that the group mean d' parameters are above zero in Experiment 3b, it is clear from the model selection procedure that the strongest account of the data requires flexibility in the participant-level d' parameters.

It is also interesting to note that while source memory for unrecognized item is lower in Experiment 3b than 3a, the estimates of source memory for recognized items are actually higher. This finding contradicts multivariate SDT models as they predict a positive correlation between item and source memory. However, the finding is consistent with dual-recollection theory (Brainerd et al., 2014) and with the experiments of Chen et al. (2018), that similarly revealed higher source performance for unrecognized items when item memory was lower.

Figure 11

Mean Source Discriminability for Unrecognized and Recognized Items in Experiments 3a and 3b

Note. Mean source discriminability (d') for unrecognized and recognized items in Experiments 3a and 3b as a function of word frequency (HF, LF). Posterior distributions of the group means of the d' parameters from the winning model ($d' \sim \text{WF}$) are represented by 95% HDIs.

With regards to Experiment 3a, there is sufficient evidence from the Bayesian statistical analysis, ROC curves, WAIC analysis, and posterior distributions from the preferred model to suggest that source memory is above chance for unrecognized items, despite the effect appearing quite subtle. For Experiment 3b, while the initial Bayesian statistical analysis yielded inconsistent findings (moderate effect for LF words, inconclusive for HF words), the subsequent WAIC analysis strongly favoured the models with non-zero discriminability for unrecognized items over the $d' = 0$ model, indicating that source memory exists for unrecognized items. While the posterior distributions of the group mean d' parameters from the preferred model included zero in the 95% HDIs, the $d' = 0$ model was still rejected as a number of participants benefited from having $d' > 0$. Nonetheless, the fact that the group mean d' parameters in Experiment 3b were not above zero suggests that the effect is more subtle than previously imagined and is difficult to capture consistently. The subtlety of the effect may be particularly problematic in the simultaneous condition, given that participants may be inclined to provide source guesses to items they don't recognize immediately, thus biasing estimates of source d' towards zero. Unfortunately, occurrences of this phenomenon are difficult to determine in the simultaneous condition as participants can anticipate the source prompt based on the preceding recognition prompt, causing both effortful and guess responses to exhibit similarly fast RTs. All in all, considering the multitude of evidence in Experiment 3a that suggests above-chance source memory for unrecognized items, along with the fact that participants may be biased towards guess responses in the simultaneous condition, we can conclude from these analyses that the effect is not just attributable to methodological confounds (as in the two blocked conditions), but is

indeed a legitimate phenomenon.

Re-Analysis of Malejka and Bröder (2016)

Perhaps the key point to take away from the current set of experiments is that while source memory in the absence of item memory appears to exist, it is a highly subtle effect at best. Even with the appropriate design (i.e., simultaneous), a large number of trials per participant, a large sample size, and a host of sophisticated analysis techniques (e.g., hierarchical Bayesian SDT models), it is difficult to establish consistently strong evidence in favour of the effect. In light of these insights, it is worth re-investigating prior studies to determine why they may not have observed an effect. The dataset of Malejka and Bröder (2016) is appropriate for such a reanalysis as it includes data for each individual participant, thus allowing for both an analysis of point estimates and a hierarchical Bayesian analysis. In their study, Malejka and Bröder (2016) observed chance-level source accuracy for unrecognized items in conditions involving the *unconditional* procedure (where source judgements were required for all items, including lures), regardless of whether participants were told that 25% (conservative bias) or 75% (liberal bias) of test items would be targets ($N = 20$ in each condition). A possibility as to why their study revealed a negative effect is that the number of presented targets was few and the recognition hit rates were high, thus resulting in a low number of unrecognized items per participant. For instance, participants in the conservative-unconditional condition studied 20 targets and exhibited a mean HR of .82 ($SE = .03$), leading to a median of just 3.50 unrecognized items per participant (and only three participants who exhibited more than five items). While participants in the liberal-unconditional condition studied 70 targets and exhibited a similar mean HR of .86 ($SE = .03$), there was still only a median of 5.00 unrecognized items per participant (and just nine participants who exhibited more than five items). As shown in Figure 5, having few unrecognized items per participant can be problematic as measures of d' for those items

become highly uncertain. It is therefore possible that the conclusions of Malejka and Bröder (2016) are the result of sample noise and not a legitimate lack of source memory for unrecognized items. To investigate this issue, we will perform an analysis of their data using hierarchical Bayesian modelling.

Two different Bayesian SDT models were devised to investigate the source data for unrecognized items and recognized items in Malejka and Bröder (2016). Both models included a criterion (c) for each response bias condition (conservative and liberal), but differed with regards to their assumptions about the discrimination (d') parameter. The first model fixed the d' parameter to zero (the $d' = 0$ model), and the second model allowed d' to vary across the two bias conditions ($d' \sim B$). The hierarchical Bayesian modelling procedure (including Stan estimation) was identical to that employed in our analysis of the present experiments (refer to Appendix B for model convergence details). Table 7 reports the ΔWAIC and $w(\text{WAIC})$ values of each model for both unrecognized and recognized items. While the $d' = 0$ model is strongly rejected for recognized items, it is the favoured model for unrecognized items. As such, there is little to no evidence from the model selection procedure to suggest that there is source memory for unrecognized items in the dataset.

Table 7

ΔWAIC and $w(\text{WAIC})$ Values for Each Hierarchical Bayesian SDT Model Applied to the Unrecognized and Recognized Items from Malejka and Bröder (2016)

Model	N	Unrecognized		Recognized	
		ΔWAIC	$w(\text{WAIC})$	ΔWAIC	$w(\text{WAIC})$
$d' = 0$	2	0.00	0.91	576.93	0.00
$d' \sim B$	4	4.66	0.09	0.00	1.00

Note. Winning models are highlighted in bold. N = number of parameters per participant.

To further investigate how the SDT models measured source discriminability for

unrecognized items, we can observe the posterior distributions of the group means of the d' parameters from the $d' \sim B$ model. Figure 12 displays the 95% HDIs of these posterior distributions alongside the source data for unrecognized and recognized items in Malejka and Bröder (2016). As can be seen, while the means of the posterior distributions align closely with the means from the data, the 95% HDIs are very wide, indicating that source d' for unrecognized items could range from anywhere between well-above or well-below chance. This suggests that the amount of available data is insufficient to produce robust estimates of the effect, even with the assistance of a hierarchical Bayesian model. As such, it could be argued that the appropriate conclusion to take from Malejka and Bröder (2016) is not that source memory is only possible for recognized items, but rather that they did not garner enough evidence to suggest otherwise.

That said, it is necessary to emphasize that the criticism of few items being presented is not unique to Malejka and Bröder (2016), nor to studies reporting no source memory for unrecognized study in general. For instance, while the experiments of Starns et al. (2008) yielded the opposite results, they share many procedural details with the aforementioned study. Instead, the issue of conclusions being drawn from insufficient data is something that appears to apply to the research area in general.

Figure 12

Mean Source Discriminability for Unrecognized and Recognized Items in Malejka and Bröder (2016)

Note. Mean source discriminability (d') for unrecognized and recognized items in Malejka and Bröder (2016) as a function of response bias condition (conservative, liberal). Posterior distributions of the group means of the SDT model's d' parameters are represented by 95% HDIs.

General Discussion

The primary aim of the present investigation was to determine whether it is possible to perform source retrieval in the absence of item memory. This aim was largely motivated by inconsistencies in the predictions of source memory models and in the empirical literature regarding the possibility of the phenomenon in question. For instance, while high-threshold models (e.g., Batchelder & Riefer, 1990) predict that source retrieval is only possible for recognized items, multivariate signal detection models (e.g., Hilford et al., 2002) predict that source retrieval is possible for unrecognized items as well (albeit with lower accuracy). Furthermore, while empirical studies that evaluate recognition and source memory on an item-by-item basis have revealed chance-level source accuracy for unrecognized items (simultaneous design; e.g., Malejka & Bröder, 2016), empirical studies that evaluate

recognition and source memory in separate blocks (blocked design; e.g., Starns et al., 2008) have revealed accuracy that is slightly above chance.

Summary of Experiments and Results

Experiments 1a and 1b. In an attempt to resolve the inconsistency in the empirical literature, we first sought to verify whether source memory for unrecognized items is indeed dependent on experimental design. To do so, two novel experiments were devised that manipulated the order of recognition and source tests within-subjects: simultaneous, blocked, and reverse blocked (where the source tests precede the recognition tests). Response modality was also manipulated between the two experiments, such that Experiment 1a required 6-point confidence ratings and Experiment 1b required yes-no judgements. Analyses of the source data using Bayesian statistical methods and hierarchical Bayesian signal detection models garnered evidence in favour of source memory for unrecognized items under each of the three design types, with the strongest and most consistent evidence being established in the two blocked conditions. These findings challenged the notion that source memory for unrecognized items is purely dependent on experimental design as the effect was also observed in the simultaneous design. The fact that the bulk of evidence in favour of the effect was established by the two blocked designs was concerning, however, as participants' item memory may not have been consistent between blocks.

Experiments 2a and 2b. To validate whether the two blocked designs were appropriate to analyse the relationship between item and source memory, two additional experiments were devised. The experiments were largely identical to the first set, except that the source tasks were replaced with remember-forget tasks, whereby participants would indicate for each target whether or not they remembered seeing it on the study list. The expectation was that the majority of recognized items in one block would be “remembered” in the alternate block, and that the majority of unrecognized items would be “forgotten” in the

alternate block. In actuality, participants frequently flipped their responses for items between blocks, placing the assumption that item memory is consistent between blocks under serious scrutiny. As such, it was concluded that the blocked designs were unsuited to exploring source memory for unrecognized items, and that evidence of the effect in such designs is likely artifactual.

Experiments 3a and 3b. Given the problems with blocked designs that were highlighted in Experiments 2a and 2b, and the fact that there was high uncertainty in the participant-level data of Experiments 1a and 1b, two new experiments were devised to establish highly robust estimates of source memory for unrecognized items in the simultaneous condition. These experiments involved roughly three times the amount of data being collected per participant in the simultaneous condition and roughly three times the sample size than in Experiments 1a and 1b. For both Experiments 3a and 3b, the model selection procedure provided conclusive evidence that source discriminability for unrecognized items was above zero. While the group-level estimates of source discriminability for unrecognized items were at chance-level for Experiment 3b, the $d' > 0$ models were still favoured as they provided a more parsimonious account of the subject-level data. Put together, the results from the simultaneous condition in Experiments 1a, 1b, 3a, and 3b suggest that source memory does exist for unrecognized items, but that it is a subtle effect at best.

Implications

The finding of source memory for unrecognized items is consistent with a number of prior studies (Ceci et al., 2010; Chen et al., 2018; Osth et al., 2018; Starns et al., 2008; Starns et al., 2013). The present investigation expands upon these works, however, by demonstrating that source memory in the absence of recognition can occur in a simultaneous design, and not solely in blocked designs which are likely to produce artifactual results. Furthermore, the

present investigation establishes the effect without the assistance of a conservative bias (Starns et al., 2008), strengthened study items (Starns et al., 2013), or context-consistent targets (Chen et al., 2018). That said, the present investigation has made it clear that source memory for unrecognized items is more difficult to measure than previously believed. Even with the appropriate experimental design, large amounts of data, and sophisticated analysis techniques, the effect appears to be highly subtle.

The present findings provide support to source memory models incorporating multivariate signal detection theory (e.g., Hilford et al., 2002) as those models allow item and source strengths to vary independently, and thus allow for the possibility of events to be recollected without being remembered. The findings also support dual-process models (e.g., Yonelinas, 1999) that incorporate a unitized familiarity mechanism, as they allow familiarity processes (governed by SDT) to contribute to source memory performance (Diana et al., 2008). In contrast, the present finding is at odds with existing source memory models incorporating high-thresholds (e.g., Batchelder & Riefer, 1990) as these models typically make explicit the assumption that source retrieval can only occur when an item is recognized. Note, however, that there is nothing inherent about the assumption of high-thresholds that precludes the ability of source memory to occur without item memory; existing high-threshold models could simply be amended to allow such a thing to occur, much in the same way that the multivariate SDT model of Hautus et al. (2008) incorporated the post-hoc assumption that unrecognized items incur source guesses. For example, the high-thresholds for item recognition and source memory in the two-high threshold source-memory model (Bayen et al., 1996) could be exchanged, such that source retrieval precedes item retrieval. In such a model, successful source retrieval could be followed by unsuccessful item retrieval, leading to source memory for unrecognized items.

Although the notion of source memory for unrecognized items remains highly

unintuitive, parallels of the effect can be made with Tulving's work on the recognition failure of recallable words (Tulving, 1974; Tulving & Wiseman, 1975). Tulving and others found that whenever recognition performance is imperfect, a small proportion of study items that can be successfully recalled during a recall test will be completely unrecognized during a recognition test. Given that recognition involves richer retrieval cues than recall (Gillund & Shiffrin, 1984), it should follow that any item that is recallable will also be recognized. However, if study items are specifically encoded in terms of their relation to other study items (e.g., as pairs), then the act of recalling other study items may provide a better cue for the target item than the item in isolation. Likewise, it is possible that if a study item were specifically encoded in conjunction with its source, then it may be easier to discriminate between particular item-source pairings (as in a source memory task), than to recognize the item in isolation, thus leading to the occurrence of source memory for unrecognized items.

Interestingly, the **hypothesis** we proposed at the outset of the present investigation – that simultaneous designs demotivate participants from performing source retrieval following a negative recognition judgement – was not supported as above-chance source accuracy for unrecognized items was observed in the condition. This observation contrasted with the results of numerous studies (Bell et al., 2017; Malejka & Bröder, 2016; Malejka et al., 2018; Onyper et al., 2010; Slotnick & Dodson, 2005; Yonelinas, 1999). Although it is difficult to diagnose the exact cause of this discrepancy, one candidate explanation is that participants throughout many of the aforementioned studies (Malejka & Bröder, 2016; Onyper et al., 2010; Slotnick & Dodson, 2005) were informed pre-test that they would be required to provide source responses to lures, and that this may have encouraged guess responses (Brainerd et al., 2014). One study went as far to instruct participants to respond in the middle of the confidence scale (indicating a preference for neither source) if they were certain a test item was a lure (Slotnick & Dodson, 2005). In contrast, participants in the present

investigation were not informed that they would have to provide source responses to lures in the simultaneous design, and as such, were not instructed how to respond in these scenarios.

Regarding the three design types we employed in Experiments 1a and 1b (i.e., blocked, reverse blocked, simultaneous), it was intriguing to observe that source performance (both for recognized and unrecognized items) was consistently the highest in the reverse blocked condition. While this finding contrasts with Chen et al. (2018, Experiment 1) who observed similar source performance between blocked and reverse blocked conditions, it makes sense intuitively. When a source test is presented in a blocked/simultaneous design, participants would have observed each item not just at study, but also during the recognition test where it was detached from its original source context. This new source context that the items are presented in would produce additional interference at retrieval (Dennis & Humphreys, 2001; Osth & Dennis, 2015), and would therefore reduce source strength for those items (Kim et al., 2012). In contrast, when a source test is presented in a reverse blocked design, participants would have only observed the test item in its original source, and thus there should be no reduction in source strength. This notion is supported by the findings of Starns et al. (2013), who observed chance-level source accuracy for unrecognized items when they were presented once at study, but above-chance accuracy when they were presented multiple times during study (i.e., source strength was increased). The implication of this theory is that source memory may indeed be possible for unrecognized items, but that chance accuracy may be observed if source strength is not high enough. Perhaps each of the experiments utilizing a simultaneous design (see Table 1) failed to detect the effect due to the fact that recognition tests precede source tests in such a design. This would be especially problematic if other aspects of the experiment also resulted in a decrease in overall source strength (e.g., weak source manipulations, lack of an encoding task, low study times, very long study lists etc...).

Finally, it is important to highlight the difference in conclusions that can be drawn from our analyses of point estimates and from our hierarchical Bayesian analysis. For instance, in the first set of experiments, our analysis of point estimates could only establish source memory for unrecognized items in the simultaneous condition for Experiment 1b. In contrast, the group-level posterior distributions of the winning model's parameters suggested that the effect was above chance in both experiments. Together with the work of Osth et al. (2018) – who observed no effect in their analysis of point estimates and a substantial effect in their hierarchical Bayesian analysis – this reveals the limitation of using non-hierarchical methods (e.g., null hypothesis significance testing, or NHST) to investigate source memory for unrecognized items, particularly when there is a small sample and/or low amount of data per participant as is often the case in similar experiments (e.g., Malejka & Bröder, 2016). Despite the fact that estimates of the effect are highly uncertain for high-performing participants with few unrecognized items, non-hierarchical methods treat each participants' mean with the same weight and have no mechanism to account for outliers, thus potentially leading to inaccurate conclusions. Hierarchical Bayesian model fitting, on the other hand, allows uncertain parameters at the participant-level to be devalued when estimating group-level parameters, and to also be adjusted in light of these group-level parameters, thus providing more sound conclusions (Boehm et al., 2018). It is possible that studies reporting no source accuracy for unrecognized items in the simultaneous condition simply used inaccurate methods (e.g., NHST, analysis of ROC curves) and/or datasets that were too small (e.g., Malejka & Bröder, 2016) to draw their conclusions.

High Source Accuracy for High Confidence Rejections

Perhaps the most intriguing finding of the present investigation was that source memory for unrecognized items was consistently the highest for items rejected with high confidence. While we initially considered the possibility that this finding was coincidental

(perhaps due to the rarity of high confidence new responses to old items), this notion was quickly dispelled as robust evidence for the phenomenon was established in both Experiments 1a and 3a, as well as in the data of several prior investigations (Onyper et al., 2010; Osth et al., 2018; Starns et al., 2013; Yonelinas, 1999). Furthermore, this trend did not appear to be due to response switching (i.e., participants accidentally responding “definitely NEW” instead of “definitely OLD”) as recognition performance was very high at the endpoints of the confidence scale. This finding is particularly intriguing as it contradicts the predictions of virtually all source memory models. It challenges models with high thresholds, such as the multinomial models of Batchelder and Riefer (1990) and the dual-process model (Yonelinas, 1999), as they predict that rejected items yield source guesses, and thus chance-level source accuracy for unrecognized items. Furthermore, it challenges multivariate SDT models (e.g., Hilford et al., 2002), as although they allow for source memory for unrecognized items, they assume that item and source strengths are correlated, and thus predict that source memory for unrecognized items should occur most often for items with moderate item strength (i.e., not items rejected with high confidence).

While it may be that this finding reveals a lack of validity for current source memory models, it is also possible that it arose due to peculiarities in participant responding. For instance, it is possible that some of the confidently rejected targets in our experiments had a high degree of memory strength but were accompanied by falsely remembered details, and these details were used as definitive evidence that the item was not seen on the study list. For example, a participant may have been presented with the word CATS in the bottom-left source during study, but they encoded the presentation as the word CAT in the bottom-left source. At test, the participant could be probed with CATS, but confidently reject the item as they have a strong memory for CAT. Nonetheless, the memory strength may be high enough such that the participant accurately attributes the item to the bottom-left source on the

subsequent source trial. If such a scenario were common, then it is possible that some of the source accuracy we observed for unrecognized items was not actually for unrecognized items, but rather for recognized items that were imperfectly encoded (and thus confidently rejected). It could also be the case the effect arose due to a disconnect between participants' memory strength and their usage of the confidence scale. Note that the majority of recognition responses were made on the endpoints of the confidence scale (e.g., 7398 “maybe NEW” and 20674 “definitely NEW” responses in Exp. 3a). It is possible that participants have a modest degree of source memory for unrecognized items, but that it is mostly attributed to the high confidence recognition responses due to a lack of usage of the medium and low confidence responses. That said, while both solutions are compelling, they do not make it clear exactly why source memory should be *higher* for high confidence rejections than low or medium confidence rejections. Ultimately, we do not fully commit to the notion that high source memory for high confidence rejections arises due to peculiarities in participant responding, but the fact that the finding strongly challenges existing source memory models means that an understanding of it constitutes an important topic for future research.

Regarding other criticisms of the present investigation, it may be questioned why we included a simultaneous design and a reverse blocked design, but not a *reverse simultaneous* design in our experiments. Given the concerns that motivation and source strength may have been impaired in the simultaneous design (as a result of recognition tests preceding source tests), it seems logical to explore the reverse simultaneous design that would arguably be free of these issues. Put simply, we consider it unlikely that this design would yield practical data as there would be a heavy bias to respond to every recognition trial with “OLD”, given the same item was just observed in the preceding source trial. Chen et al. (2018, Experiment 2) were able to utilize a similar design by separating source and item trials for the same item

with trials for other items, but doing so also increases the likelihood that item memory would change between trials for the same item, potentially leading to artifactual results (c.f., blocked designs).

Conclusion

Empirical data from three sets of experiments and subsequent fits from hierarchical Bayesian models suggest that source memory for unrecognized items is indeed possible, but that it is a highly subtle effect and is more difficult to measure than previously believed. The experiments demonstrated: (a) that the effect can occur when recognition and source memory are tested simultaneously, despite prior studies suggesting otherwise; and (b) that while the effect can also occur when recognition and source memory are tested in separate blocks, such results are likely artifactual, resulting from frequent changes in item memory between blocks. The overall findings support existing models of item and source memory that incorporate continuous latent strengths, as opposed to all-or-none processes or high-thresholds. It is suggested that findings in prior literature of there being no source memory for unrecognized items may be attributable to design elements that decrease overall source strengths and/or encourage guess responses, small sample sizes and/or low amounts of data per participant, and to a reliance on non-hierarchical analysis methods that are insensitive to uncertainties in participant-level data that may arise when source performance is conditioned on unrecognized items.

References

- Annis, J., Miller, B. J., & Palmeri, T. J. (2017). Bayesian inference with Stan: A tutorial on adding custom distributions. *Behavior Research Methods*, 49(3), 863-886.
- Baayen, R. H., Piepenbrock, R., & Gulikers, L. (1996). The CELEX Lexical Database (CD-ROM). *Linguistic Data Consortium*.
- Banks, W. P. (2000). Recognition and source memory as multivariate decision processes. *Psychological Science*, 11(4), 267-273.
- Batchelder, W. H., & Riefer, D. M. (1990). Multinomial processing models of source monitoring. *Psychological Review*, 97(4), 548-564.
- Bayen, U. J., Murnane, K., & Erdfelder, E. (1996). Source discrimination, item detection, and multinomial models of source monitoring. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(1), 197-215.
- Bell, R., Mieth, L., & Buchner, A. (2017). Emotional memory: No source memory without old-new recognition. *Emotion*, 17(1), 120-130.
- Benjamin, A. S., Diaz, M., & Wee, S. (2009). Signal detection with criterion noise: Applications to recognition memory. *Psychological Review*, 116(1), 84-115.
- Benjamin, A. S., Tullis, J. G., & Lee, J. H. (2013). Criterion noise in ratings-based recognition: Evidence from the effects of response scale length on recognition accuracy. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(5), 1601-1608.
- Boehm, U., Marsman, M., Matzke, D., & Wagenmakers, E. J. (2018). On the importance of avoiding shortcuts in applying cognitive models to hierarchical data. *Behavior Research Methods*, 50(4), 1614-1631.
- Brainerd, C. J., Gomes, C. F. A., & Moran, R. (2014). The two recollections. *Psychological Review*, 121(4), 563-599.

- Brainerd, C. J., Reyna, V. F., Holliday, R. E., & Nakamura, K. (2012). Overdistribution in source memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(2), 413-439.
- Ceci, S. J., Fitneva, S. A., & Williams, W. M. (2010). Representational constraints on the development of memory and metamemory: A developmental–representational theory. *Psychological Review*, 117(2), 464-495.
- Chen, X. R., Gomes, C. F. A., & Brainerd, C. J. (2018). Explaining recollection without remembering. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(12), 1921-1930.
- Cook, G. I., Marsh, R. L., & Hicks, J. L. (2006). Source memory in the absence of successful cued recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(4), 828-835.
- deBettencourt, M. T., Norman, K. A., & Turk-Browne, N. B. (2018). Forgetting from lapses of sustained attention. *Psychonomic Bulletin & Review*, 25(2), 605-611.
- DeCarlo, L. T. (2003). An application of signal detection theory with finite mixture distributions to source discrimination. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 767-778.
- Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2008). The effects of unitization on familiarity-based source memory: testing a behavioral prediction derived from neuroimaging data. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(4), 730-740.
- Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, 62(3), 145.
- Fox, J., & Osth, A. F. (2021, June 17). Does source memory exist for unrecognized items? Retrieved from osf.io/my93p

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013).

Bayesian data analysis. CRC press.

Gelman, A., Lee, D., & Guo, J. (2015). Stan: A probabilistic programming language for

Bayesian inference and optimization. *Journal of Educational and Behavioral Statistics*, 40(5), 530-543.

Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple

sequences. *Statistical Science*, 7(4), 457-472.

Gillund, G., & Shiffrin, R. M. (1984). A Retrieval Model for Both Recognition and Recall.

Psychological Review, 91(1), 1-67.

Glanzer, M., & Adams, J. K. (1985). The mirror effect in recognition memory. *Memory &*

Cognition, 13(1), 8-20.

Glanzer, M., Hilford, A., & Kim, K. (2004). Six Regularities of Source Recognition. *Journal*

of Experimental Psychology: Learning, Memory, and Cognition, 30(6), 1176-1195.

Hautus, M. J., Macmillan, N. A., & Rotello, C. B. (2008). Toward a complete decision model

of item and source recognition. *Psychonomic Bulletin & Review*, 15(5), 889-905.

Hilford, A., Glanzer, M., Kim, K., & DeCarlo, L. T. (2002). Regularities of source

recognition: ROC analysis. *Journal of Experimental Psychology: General*, 131(4), 494-510.

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context.

Journal of Mathematical Psychology, 46(3), 269-299.

Jeffreys, H. (1961). *Theory of probability*. Oxford University Press.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source Monitoring. *Psychological*

Review, 114(1), 3-28.

Kass, R. E., & Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical*

Association, 90(430), 773-795.

- Kellen, D., Erdfelder, E., Malmberg, K. J., Dubé, 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., & Klauer, K. C. (2014). Modeling source-memory overdistribution. *Journal of Memory and Language*, 76, 216-236.
- Kim, K., Yi, D. J., Raye, C. L., & Johnson, M. K. (2012). Negative effects of item repetition on source memory. *Memory & Cognition*, 40(6), 889-901.
- Klauer, K. C., & Kellen, D. (2010). Toward a complete decision model of item and source recognition: A discrete-state approach. *Psychonomic Bulletin & Review*, 17(4), 465-478.
- Koop, G. J., Criss, A. H., & Malmberg, K. J. (2015). The role of mnemonic processes in pure-target and pure-foil recognition memory. *Psychonomic Bulletin & Review*, 22(2), 509-516.
- Kurilla, B. P., & Westerman, D. L. (2010). Source memory for unidentified stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(2), 398-410.
- Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.
- Luce, R. D. (1963). A threshold theory for simple detection experiments. *Psychological Review*, 70(1), 61-79.
- Macmillan, N. A., & Creelman, C. D. (2004). *Detection theory: A user's guide*. Psychology Press.
- Malejka, S., & Bröder, A. (2016). No source memory for unrecognized items when implicit feedback is avoided. *Memory & Cognition*, 44(1), 63-72.
- Malejka, S., Ciecior, A. S., Möller, C., & Bröder, A. (2018, November). *No source memory without item memory despite increased motivation for source judgements after "new"*

- responses*. Poster presented at the 58 th Annual Meeting of the Psychonomic Society (Psychonomics), Vancouver, BC, Canada.
- Mulligan, N. W., & Osborn, K. (2009). The modality-match effect in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(2), 564-571.
- Onyper, S. V., Zhang, Y. X., & Howard, M. W. (2010). Some-or-none recollection: Evidence from item and source memory. *Journal of Experimental Psychology: General*, 139(2), 341-364.
- Osth, A. F., Fox, J., McKague, M., Heathcote, A., & Dennis, S. (2018). The list strength effect in source memory: Data and a global matching model. *Journal of Memory and Language*, 103, 91-113.
- Peirce, J. W., Gray, J. R., Simpson, S., MacAskill, M. R., Höchenberger, R., Sogo, H., Kastman, E., Lindeløv, J. (2019). PsychoPy2: experiments in behavior made easy. *Behavior Research Methods*. 10.3758/s13428-018-01193-y
- Plummer, M. (2003, March). 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, No. 125.10, pp. 1-10).
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56(5), 356-374.
- Slotnick, S. D., & Dodson, C. S. (2005). Support for a continuous (single-process) model of recognition memory and source memory. *Memory & Cognition*, 33(1), 151-170.
- Starns, J. J., Hicks, J. L., Brown, N. L., & Martin, B. A. (2008). Source memory for unrecognized items: Predictions from multivariate signal detection theory. *Memory & Cognition*, 36(1), 1-8.

- Starns, J. J., & Ksander, J. C. (2016). Item strength influences source confidence and alters source memory zROC slopes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(3), 351-365.
- 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(4), 527-539.
- Starns, J. J., Pazzaglia, A. M., Rotello, C. M., Hautus, M. J., & Macmillan, N. A. (2013). Unequal-strength source zROC slopes reflect criteria placement and not (necessarily) memory processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(5), 1377-1392.
- Stefan, A. M., Gronau, Q. F., Schönbrodt, F. D., & Wagenmakers, E.-J. (2019). A tutorial on Bayes Factor Design Analysis using an informed prior. *Behavior Research Methods*, 51(3), 1042-1058.
- Tulving, E. (1974). Cue-dependent forgetting: When we forget something we once knew, it does not necessarily mean that the memory trace has been lost; it may only be inaccessible. *American Scientist*, 62(1), 74-82.
- Tulving, E., & Wiseman, S. (1975). Relation between recognition and recognition failure of recallable words. *Bulletin of the Psychonomic Society*, 6(1), 79-82.
- Wagenmakers, E. J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, 14(5), 779-804.
- Wagenmakers, E. J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology*, 21(5), 641-671.
- Wagenmakers, E. J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11(1), 192-196.

- Wagenmakers, E. J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., ... & Morey, R. D. (2018). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic Bulletin & Review*, 25(1), 58-76.
- Wallace, W. P. (1982). Distractor-free recognition tests of memory. *The American Journal of Psychology*, 95, 421–440.
- Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11(Dec), 3571-3594.
- Yonelinas, A. P. (1999). The contribution of recollection and familiarity to recognition and source-memory judgements: A formal dual-process model and an analysis of receiver operating characteristics. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(6), 1415-1434.
- Zhou, J., Osth, A. F., Lilburn, S. D., & Smith, P. L. (2021). A circular diffusion model of continuous-outcome source memory retrieval: Contrasting continuous and threshold accounts. *Psychonomic Bulletin & Review*, 28, 1112-1130.

Appendix A

Prior Distributions on Model Parameters

In our hierarchical Bayesian analyses of the present experiments and the Malejka and Bröder (2016) data, all participant-level parameters were sampled from group-level distributions with mean μ and standard deviation σ . For each of the experiments, source criteria and d' parameters were sampled from normal distributions using a log scale:

$$\begin{aligned} d' &\sim \text{Normal}(\mu_d, \sigma_d) \\ c &\sim \text{Normal}(\mu_c, \sigma_c) \\ c_3 &\sim \text{Normal}(\mu_{c3}, \sigma_{c3}) \end{aligned} \tag{A1}$$

where c_3 is the central criterion of the five criteria used for Experiments 1a and 3a. The four remaining criteria for Experiments 1a and 3a were calculated relative to the central criterion as:

$$\begin{aligned} c_1 &= c_3 - \phi_2 - \phi_1 \\ c_2 &= c_3 - \phi_2 \\ c_4 &= c_3 + \phi_4 \\ c_5 &= c_3 + \phi_4 + \phi_5 \end{aligned} \tag{A2}$$

where the lower numbers (1 and 2) represent the liberal criteria and the higher numbers (4 and 5) represent the conservative criteria. This form of parameterization was selected as it guaranteed a partial ordering of the response criteria. Each of the ϕ parameters were sampled from truncated normal (TN) distributions that were truncated from zero to infinity:

$$\begin{aligned} \phi_1 &\sim TN(\mu_{\phi1}, \sigma_{\phi1}, 0, \infty) \\ \phi_2 &\sim TN(\mu_{\phi2}, \sigma_{\phi2}, 0, \infty) \\ \phi_4 &\sim TN(\mu_{\phi4}, \sigma_{\phi4}, 0, \infty) \\ \phi_5 &\sim TN(\mu_{\phi5}, \sigma_{\phi5}, 0, \infty) \end{aligned} \tag{A3}$$

The priors used for each of the group-level parameter distributions were relatively non-informative:

$$\mu_d \sim \text{Normal}(0,10)$$

$$\mu_c \sim \text{Normal}(0,1)$$

$$\mu_\phi \sim \text{TN}(.5, .5, 0, \infty)$$

$$\sigma_{d,c,\phi} \sim \text{Gamma}(1,1) \tag{A4}$$

Appendix B

Model Convergence

Experiments 1a and 1b

For each of the model fits to the Experiment 1 data, there was a high degree of converge for the MCMC chains. Gelman-Rubin convergence diagnostics (Gelman & Rubin, 1992) were lower than 1.1 for each parameter in each of these models. For Experiment 1a, the preferred model fits exhibited maximum Gelman-Rubin statistics of 1.001 ($M < 1.001$) and 1.034 ($M = 1.003$) alongside mean effective sample sizes (ESS; Gelman et al., 2013) of 41983 and 7083 for recognized and unrecognized items respectively. For Experiment 1b, the preferred model fits exhibited maximum Gelman-Rubin statistics of 1.002 ($M < 1.001$) and 1.003 ($M < 1.001$) alongside mean ESS values of 16890 and 9273 for recognized and unrecognized items respectively.

Experiments 3a and 3b

Gelman-Rubin convergence diagnostics were lower than 1.1 for each parameter in each of models fit to the Experiment 3 data. For Experiment 3a, the preferred model fits exhibited maximum Gelman-Rubin statistics of 1.0003 ($M < 1.001$) and 1.022 ($M = 1.002$) alongside mean ESS values of 42008 and 7654 for recognized and unrecognized items respectively. For Experiment 3b, the preferred model fits exhibited maximum Gelman-Rubin statistics of 1.001 ($M < 1.001$) and 1.001 ($M < 1.001$) alongside mean ESS values of 17985 and 13122 for recognized and unrecognized items respectively.

Malejka and Bröder (2016)

Gelman-Rubin convergence diagnostics were lower than 1.1 for each parameter in each of the models fit to the Malejka and Bröder (2016) data. The preferred model fits exhibited maximum Gelman-Rubin statistics of 1.001 ($M < 1.001$) and 1.001 ($M = 1.001$)

alongside mean ESS values of 11176 and 5770 for recognized and unrecognized items respectively.

Appendix C

Posterior Checks for Preferred Models

Figures C1, C2 and C3 present the Experiment 1a ROC data for recognized and unrecognized items across the blocked, reverse blocked and simultaneous conditions respectively, alongside simulations from the preferred model ($d' \sim D$, WF). Figure C4 presents the Experiment 3a ROC data for recognized and unrecognized items alongside simulations from the same model. As can be seen, there is a high degree of consistency between the ROC data and the model simulations across both experiments, indicating that the model describes the major aspects of the data.

Figure C1

Experiment 1a ROC Data for Recognized and Unrecognized Items in the Blocked Condition

Note. ROCs are presented for both HF words (left panel) and LF words (right panel). Black markers represent simulated ROC data from the preferred model ($d' \sim D$, WF).

Figure C2

Experiment 1a ROC Data for Recognized and Unrecognized Items in the Reverse Blocked

Condition

Note. ROCs are presented for both HF words (left panel) and LF words (right panel). Black markers represent simulated ROC data from the preferred model ($d' \sim D$, WF).

Figure C3

Experiment 1a ROC Data for Recognized and Unrecognized Items in the Simultaneous Condition

Note. ROCs are presented for both HF words (left panel) and LF words (right panel). Black markers represent simulated ROC data from the preferred model ($d' \sim D$, WF).

Figure C4*Experiment 3a ROC Data for Recognized and Unrecognized Items*

Note. ROCs are presented for both HF words (left panel) and LF words (right panel). Black markers represent simulated ROC data from the preferred model ($d' \sim D$, WF).