Neuropsychologia 90 (2016) 94–109

Contents lists available at ScienceDirect

Neuropsychologia

journal homepage: www.elsevier.com/locate/neuropsychologia

The unitary zROC slope in amnesics does not reflect the absence of recollection: critical simulations in healthy participants of the zROC slope

Chen Didi-Barnea, Ziv Peremen, Yonatan Goshen-Gottstein *

School of Psychological Sciences, Tel Aviv University, Ramat Aviv, 69978, Israel

ARTICLE INFO

Article history:
Received 22 December 2015
Received in revised form
10 April 2016
Accepted 24 May 2016
Available online 25 May 2016

Keywords:
Amnesia
Recognition
Remember-know
Strength
Confidence
zROC slope

ABSTRACT

The functional deficit underlying amnesia has been argued to be in recollective processing. This idea is based on the DPSD model, wherein recognition comprises a mixture of recollection and familiarity signals, with familiarity conforming to an equal-variance signal-detection mechanism while recollection is binary. This model interprets the greater variance for targets than for lures revealed in sub-unit zROC slopes, to be a consequence of the mixture of two signals. Importantly, equal variance between targets and lures is found in amnesic, and is consequently interpreted by DPSD to reflect impairment to recollection alongside the sparing of familiarity. Here, we pointed to a logical fallacy in this interpretation. We then asked participants, in two experiments, to make remember-know (RK) and confidence judgments. Simulating equal variance in healthy participants, we either excluded from the analysis 'remember' responses, reflecting recollection, or the most accurate memories, reflecting strength. We found that only the exclusion of the strongest responses led to equal-variance distributions. In addition, we found that accuracy was associated with an interlaced ordering of RK response groups nested under confidence, a pattern hard to reconcile with classic recognition models (DPSD, UVSD). This pattern can, however, be accommodated by the Continuous Dual Process (CDP) model (Wixted and Mickes, 2010), wherein both familiarity and recollection are continuous signals. Amnesia may thus be characterized as the inability to form strong memories, recollection as well as familiarity.

1. Introduction

An important goal in the cognitive neuropsychology of memory is to delineate the functional loss in amnesia. Over the past 50 years, a multitude of processes have been suggested as being impaired in anterograde amnesia, beginning with the idea that it comprises a deficit in long term memory (e.g., Milner, 1966; Wickelgren, 1968; cf., Davelaar et al., 2005; Talmi and Goshen-Gottstein, 2006), and moving on to a long list of suggestions of possible deficits within long-term memory, including a deficit to declarative memory (as compared to procedural memory; Cohen and Squire, 1980; Manns et al., 2003), a deficit to episodic memory (as compared to semantic memory; Viskontas et al., 2000), a deficit in retrieval (Nadel and Moscovitch, 1997), a storage deficit (e.g., Hardt et al., 2009; Mayes, 1995), a deficit in conscious recollection as indexed by explicit tests of memory (for reviews, see Cohen and Eichenbaum, 1993; Moscovitch, 1982), a deficit in relational memory (Cohen et al., 1997; Ryan et al., 2000; but see Goshen-Gottstein et al., 2000), as an impairment of detail generation and binding (Rosenbaum et al., 2009), as an impairment in autobiographical memory (e.g., Rosenbaum et al., 2004), and—most relevant to our present concerns—as an impairment in recollective processing, sparing familiarity (Gilboa et al., 2006; Hirst et al., 1988; Schacter et al., 1984; Turriziani et al., 2008; Yonelinas et al., 1998). In this article, we argue against the idea that a recollective deficit characterizes amnesia. Instead, we suggest that amnesia may represent a deficit in the formation of strong traces, comprising both recollection and familiarity.

The investigations reported in this article are guided by the notion that a better understanding of the nature of the deficit underlying (anterograde) amnesia can be obtained if pursuant to a carefully designed manipulation, performance in neurologically intact participants can be shown to simulate the deficit observed in amnesia. Such manipulations include pharmaceutical interventions (e.g., Hardt et al., 2009), transcranial magnetic stimulations (e.g., Bolognini and Ro, 2010) and cognitive tasks (e.g., Dunbar and Sussman, 1995; Moscovitch, 1994). Experimentally-induced impairments have the benefit that they can be subject to systematic

---

* Corresponding author.
E-mail address: goshengott@gmail.com (Y. Goshen-Gottstein).

http://dx.doi.org/10.1016/j.neuropsychologia.2016.05.026
0028-3932/

investigations, with the purpose of expanding our understanding of the mechanism underlying the observed behavior.

Here, we focus on a pivotal finding regarding the performance of amnesic patients in recognition tests, which was uncovered in analyses of ROC curves. The pattern revealed in ROC analysis (described in detail below), suggests that amnesia may be characterized as impairment in the recollective process. In validation of this suggestion was the finding that the amnesic ROC pattern (Aggleton et al., 2005; Gilboa et al., 2006; Yonelinas et al., 1998) could be revealed in healthy participants, when applying a simulation procedure aimed at disabling recollective processing (Yonelinas, 2001). In this article, we demonstrate how critical investigations of this simulation procedure should change our conceptualization of the functional deficit underlying amnesia.

The suggestion that recollective processes are impaired in amnesia is primarily based on studies of recognition memory. In recognition, participants are presented with items at study, and at test are asked to distinguish between targets (studied items) and lures (unstudied items). Dual-process models of recognition (the Dual-Process Signal Detection (DPSD), Yonelinas (1994); the Variable-Recollection Dual-Process (VRDP), Onyper et al. (2010); the Continuous Dual-Process (CDP), Wixted and Mickes (2010); the Sum-difference Theory of Remembering and Knowing (STREAK), Rotello et al. (2004)) propose that performance on this task is an outcome of two qualitatively different signals, recollection and familiarity. As we shall see below, impaired processing of one of these signals—recollection—has been argued to be the source of deficit in amnesia.

The recollection signal entails memory for details of an episodic event, which may, or may not, be available to consciousness. For example, remembering not only that you met a certain person before, but also that person's shirt color, a background song, and topics of conversation in the specific encounter. The recollection process may fail to provide episodic details, yet often people still report experiencing a feeling of knowing that the event occurred (e.g., that the person has been previously encountered). This feeling of knowing corresponds to output from the second signal—familiarity. The notion that recognition comprises both recollection and familiarity has been supported in numerous studies using behavioral methodologies, fMRI, scalp electroencephalography, human neuropsychology and the study of recognition in non-humans (Aggleton et al., 2005; Curran, 2000; Diana et al., 2006; Eichenbaum et al., 2007; Fortin et al., 2004; Rugg and Curran, 2007; Rugg and Yonelinas, 2003; Woodruff et al., 2006; Yonelinas et al., 2005; Yonelinas, 2002; Yovel and Paller, 2004).

A triad of findings (see details below) turns out to be critical in constraining theories of recognition memory as well as for delineating the functional deficit underlying amnesia. The three findings were uncovered in the analysis of ROC curves. An ROC curve describes the function relating the proportion of correctly recognized target items (i.e., hit rate) to the proportion of incorrectly recognized lure items (i.e., false alarm rate) across variations in response criteria. The notion of response criteria refers to the bias in making a positive recognition response, irrespective of the amount of recognition-relevant information (e.g., familiarity and/or recollection) that is available for a particular stimulus. When the underlying lure and target distributions are normal—a standard assumption of signal-detection theory (SDT; Swets and Green, 1963)—the ROC curve plotted in z-space (zROC) is linear. Importantly, the slope of the linear zROC curve has been shown to represent the ratio of lure-to-target distribution variance (or, strictly speaking, standard deviation). Specifically, a zROC slope equal to 1 indicates equal variance between the target and lure distributions and a slope smaller than 1 indicates a larger variance of the target distribution.

The first of the triad of findings was that the zROC slope in standard tests of recognition in healthy participant is smaller than 1 ($\sim$0.8), indicating a higher variance of the target distribution than the lure distribution (e.g., Ratcliff et al., 1992). This finding is important, in that standard SDT models assume equal variance for target and lure distributions. To account for the inequality of variance, profoundly different models have been postulated, ranging from an assortment of dual-process models to models that posit only a single process. The second finding was that the slope of the zROC curve in amnesics is equal to 1, indicating equal variance for target and lure distributions, as assumed by SDT. As we shall see, this finding has been interpreted as evidence that the function evidence in amnesia is that of impairment to recollective processing (an interpretation we challenge in this article). The third finding was that of equal variance of targets and lures, when a simulation procedure was used to eliminate recollective processing in neurologically intact participants (Yonelinas, 2001). We now describe two classic models, DPSD—a dual-process model, and UVSD—a single process model. The two models represent radically different interpretations of the inequality of variance observed in healthy individuals and of the equality observed in amnesia and under the simulation procedure.

1.1. DPSD

The Dual Process Signal Detection (DPSD) model (Yonelinas, 1994), suggests that a familiarity signal is available for all items. In contrast, recollection is available for some, but not all, of the target items. Specifically, the familiarity signal, assumed to be governed by the principles of SDT, is argued to mediate both the lure and the target distributions. Here, the effect of studying the items is reflected by a uniform—constant—boost (increase) to the lure distribution. Because the boost is uniform for all studied items, the resultant target distribution is equal in variance to that of the lure distribution. Thus, the familiarity signal conforms to an equal-variance signal detection mechanism. In addition, the target distribution alone is influenced by retrieval of a binary, high threshold (Macmillan and Creelman, 1991) recollection signal that is generated only for some of the items that have been studied. Importantly, the empirically-observed inequality of variance is interpreted by DPSD to be the outcome of the equal-variance familiarity signal available for both the lure and target items, and the binary recollection signal available for only target items.

According to DPSD, if recognition were to be tested in patients with a deficit in recollective processing, then their performance should reflect the spared familiarity process, yielding a zROC slope of 1 (reflecting an equal-variance distribution). Such a unitary zROC slope was in fact observed in densely amnesic patients and patients with hippocampal injury (Aggleton et al., 2005; Yonelinas et al., 2002, 1998), suggesting an impairment of recollection as the source of the functional deficit in amnesia. Bolstering the notion of a function deficit in recollection in amnesia was the finding of the simulation procedure, which presumably mimicked the functional deficit of the recollective process in healthy participants. To this end, the Remember-know task (Gardiner, 1988; Tulving, 1985) was used in conjunction with zROC analysis. In the Remember-know (R-K) task, participants are required to provide a subjective classification of positive recognition judgments into recognition that entails 'remembering' (R) something specific about the study episode as opposed to recognition that only comprises 'knowing' (K) that an item was studied, with failure to retrieve any specific detail from the learning episode. R judgments have been argued to reflect the operation of the recollection signal, with K judgments associated with familiarity signal (with a possible correction for an underestimation of this process; Yonelinas and Jacoby, 1995). Systematic investigation of performance on the R-K task have supplied converging evidence in support of DPSD (for review, see Yonelinas, 2002; but see Dunn, 2004).

To simulate recognition-without-recollection, Yonelinas, 2001 used a two-stage procedure. The first stage comprised a 6-point confidence recognition test (ranging from 1- ‘sure new’ to 6- ‘sure old’). This was followed in second stage, for each recognized item, by an R-K judgment. The typical smaller-than-one (i.e., ~0.8) zROC slope was found. Next, Yonelinas performed a zROC analysis only for items for which K responses were given. These items can be likened to the items recognized by amnesic patients, for whom the recollective process is impaired. Critically, when R responses were selectively excluded from the zROC calculations, reserving for analysis only items judged K, the zROC slope increased to 1. This finding—which is the focus of the current article—was interpreted to suggest that participants’ R judgments reflect recollection-based recognition. Hence, when R responses were selectively excluded from the analysis, a pattern predicted by the DPSD model emerged, that of equal variance for the target and lure distribution. The zROC slope of 1 which was found after removal of R responses, therefore, strengthened the validity of DPSD model as well as the interpretation of functional deficit in amnesia as a failure of the recollective processes.

Critically, the argument from zROC analysis that amnesia (or for that matter, the simulation procedure) indicates a deficit in recollective processing entails a logical fallacy. To make the argument, it was first assumed that DPSD provides a true characterization of recognition memory. If so, the argument continued, then conditional on a recollective deficit, a unitary zROC slope should be revealed, in amnesics as well as in K responses of healthy participants. This prediction is entirely correct. However, inference in the reverse direction constitutes the fallacy of inferring the consequent and is, therefore, invalid. That is, conditional on finding a unitary slope patients (or in K responses of healthy participants), evidence for a recollective deficit cannot be inferred. Indeed, in this article, we argue that mechanisms other than familiarity—specifically, the existence of weak memories—may mediate the unitary zROC slope. Such mechanisms are described by a fundamentally different model—UVSD—a model that is based on memory strength. We now turn to describe UVSD, and its account of the triad of findings.

1.2. UVSD

The unequal variance signal detection (UVSD) model (Dunn, 2004) is premised on the notion that recognition judgments are based on single, sum-of-evidence signal—a continuum of memory strength (Donaldson, 1996; Dunn, 2004, 2008) according to the principles of signal-detection theory. Specifically, a subjective criterion is set by participants along the strength continuum, and judgments made depending on whether an item's strength exceeds (‘old’) or does not exceed (‘new’) the criterion. Like old-new judgments, R-K judgments too are interpreted by UVSD as reflecting the setting of a low and high criterion on the strength distribution. Thus, on trials in which the signal exceeds the high criterion (i.e., trials with ‘strong memory’), an R response is made, with a K response given on trials in which the signal lies between the two criteria (i.e., ‘weak memory’ trials) (see Donaldson, 1996). Finally, items are judged as ‘new’, when their signal falls below the low criterion. Thus, the more parsimonious single-process account suggests that the distinction between R and K reflects different quantitative degrees of memory strength rather than qualitatively distinct memory processes.

To account for the inequality of variance observed in ROC analysis, UVSD proposes a variable—rather than a uniform—boost to the mnemonic strength of items during the study episode. The variable increase could be a consequence of, for example, variations of attention at encoding (Ratcliff et al., 1992; Wixted, 2007). Thus, because the target distribution consists of both weak and strong memory traces, its (mean and) variability is found to be greater than that of the lure distribution—leading to the robust lower-than-one zROC slope pattern that has been systematically found in the empirical data. UVSD provides a viable alternative for many of the behavioral (Donaldson, 1996; Dunn, 2004, 2008; Rotello and Zeng, 2008; Squire et al., 2007; Wixted and Stretch, 2004) and neuroscientific (e.g., Wixted, 2009) results that have been described as supporting the DPSD model.

UVSD interprets the equal variance observed in amnesic patients, reflected by the unitary zROC slope, by suggesting that the target distribution in amnesics consists primarily of weak memory traces (Squire et al., 2007). According to this suggestion, the encoding of events by amnesics yields only a small, boost in memory strength, if any boost at all. Consequently, the target distribution (mean and its) variability would not be much greater than the lure distribution. Hence, amnesia may be characterized by the inability to bolster memory traces, rather than as a deficit specifically targeting recollection-based processes.

Support of the idea that a zROC slope of 1 may reflect weak memories—and that amnesia should therefore be characterized as the inability to form strong memories, neither recollection nor familiarity —was found in two studies. First, in healthy individuals, using experimental manipulations to enhanced memory performance (e.g., study time, encoding task, repetition)—purportedly reflecting stronger memories—(Glanzer et al., 1999) found that with the increase in performance, the zROC slope became progressively smaller than 1. Second, Wais et al. (2006) examined recognition memory in hippocampal-injured patients, who were presented with short study lists, so as to enhance amnesic performance and match it to that of controls. Importantly, the zROC slopes of the amnesics decreased to values approaching those found in controls, that is, values significantly lower than 1. Together, these two findings suggest that the slope of the zROC slope may better be characterized as indicative of memory strength, rather than as reflecting whether recognition judgments involved recollection.

Here, we further investigated the validity of the strength interpretation of amnesia by exploring, in healthy individuals, the effects of various exclusion schemes on the slope of the zROC curve. We suggested that the exclusion of R responses from the ROC analysis performed by Yonelinas, 2001, may have been confounded with removal of the strongest memory traces. Therefore, the equal variance following the selective exclusion of R responses does not necessarily indicate the lone operation of the familiarity signal, as suggested by the DPSD model. Exclusion of these responses may just as likely yield lure and target distributions comprising only weaker memories, for which the target distribution more closely resembles that of the lures, both in mean and variance.

Note that according to the strength interpretation, there is no special status to a unitary zROC slope. That is, the exclusion of the strongest responses is predicted to give rise to the zROC slope, but not necessarily a rise to 1. In contrast, according to DPSD, a rise in the slope to values that are lower than 1, would not provide support for the theory. Thus, DPSD makes a much stronger prediction than a strength interpretation, in that only DPSD predicts a slope of precisely 1.[1]

To test this idea, we examined whether the selective removal in healthy individuals of the strongest responses—rather than of recollection-bases responses—would yield an outcome of equal variance (or at least of a rise in the magnitude of the zROC slope)

---

[1] In Section 4.3, we simulated the removal of the strongest responses from a UVSD model and showed that a slope no-smaller than 1 turns up surprisingly often. Also see Footnote 3.

comparable to that observed following the selective exclusion of R responses. Indeed, under certain conditions, it may provide a better account for the increase in the zROC slope after R responses are excluded.

We asked participants to rate both the confidence of their judgments on 6-point scale as well as their subjective experience (R-K). According to DPSD, it is removal of R responses, reflecting a qualitatively unique process—recollection, which leads to an increase of the zROC slope to 1. In contrast, UVSD predicts that it is removal of a responses which are quantitatively strong that may provide a better account for the changes in the zROC slope following selective exclusion.

2. Experiment 1

2.1. Material and methods

2.1.1. Participants

Eighteen Tel-Aviv University students (2 males, 24 ± 2.3 years old) participated in the experiment for monetary compensation (approximately 7$ US) or credit in an undergraduate psychology course. One participant was excluded from the analysis, because her zROC slope was 0.38, which was more than three SDs lower than the mean slope (=0.71). Additionally, to obtain an adequate number of responses per bin, we excluded an additional participant who had less than 15% of the responses categorized as K judgments, thus, leading to two response groups, 6K and 5K, comprising less than 0.06% of trials.

2.1.2. Stimuli

The stimuli consisted of 480 Hebrew 2–3 syllables words, 3–4 letters long, divided into six lists of 80 words each. The lists had similar frequency, measured by the average number of Google© search-engine results for each of the words (M=1810 K, SD=154 K). The six lists were divided into three pairs of lists, with each pair used in a separate block in the experiment, for a total of three study-test blocks. In each block, the words of one of the lists were presented at study and the words from both lists, at test. To ensure that all words appeared at test an equal number of times as studied and as unstudied, across participants, each list appeared an equal number of times at study. The order of the blocks was randomized for each participant.

2.1.3. Procedure

Participants were informed that following each study block a test block would be administered, wherein their memory would be tested. They were told that altogether, they would be presented with three such study-test blocks. During each study phase, 80 single words were presented randomly for 2 s each, for a total of 240 study words across the experiment. Before the presentation of each word, a cross appeared in the middle of the screen for 250 ms followed by an empty screen for additional 250 ms. Participants were briefed to say aloud the displayed words.

Following each study phase, participants completed a recognition test, wherein 80 studied and 80 unstudied words were presented one at a time in random order on the screen, for a total of 480 words throughout the experiment. Participants made recognition judgments on the computer keyboard using a 6-point confidence scale. They were instructed to respond '6' if they were sure the item had been presented, '5' if they were less sure, and '4' if very unsure. They were instructed to respond '1' if they were sure the word had not been presented, '2' if they were less sure and '3' if they were very unsure. Thus, responses 1–3 correspond to a 'new' recognition judgment and 4–6, to an 'old' judgment. The confidence scale appeared on the screen with the appropriate descriptive labels, whenever confidence judgment was required.

If an 'old' judgment (i.e. confidence levels 4–6) was made, participants were asked to make an R-K judgment using the left/right arrow keys on the computer keyboard. They were asked to respond R if they could recollect specific details regarding the item's study presentation and to respond K if they felt that this word was studied, without recollecting details from the actual appearance of the word in the study list. The instructions were translated to Hebrew from those used by Rajaram (1993). Instructions included a clarification that R-K judgments are not necessarily related to the rating of confidence (but see Yonelinas (2001); General Discussion, Section 4.2). Participants were tested individually, with each session lasted approximately 45 min

Displays were generated by an Intel Pentium 4 computer attached to a 15-inch CRT monitor, using 1024 × 768 resolution graphics mode. Viewing distance was set at 70 cm. The stimuli were presented in 20-point Arial font.

2.2. Results

2.2.1. General

Across participants, mean performance yielded a 71% hit rate (HR; SE=2.9%) and a 20% false alarm rate (FAR; SE=2.2%). More specifically, R judgments reflected a 42% HR and 5% FAR, whereas K judgments were associated with a 30% HR and a 15% FAR. Across hits and false alarms, R judgments were associated with higher confidence ratings (M=5.55; SE=0.12) than K judgments (M=4.68; SE=0.13), t_{15}=3.87, p=.002. In addition, recognition accuracy of studied items was higher for R judgments (M=89%; SE=0.03) than for K judgments (M=67%; SE=0.03), t_{15}=4.11, $p < 0.001$.

Throughout the article, when computing accuracy for a particular response group (e.g., R, K; confidence level 4; confidence level 5, K5, R6), accuracy was computed as proportion correct—defined as hits/(hits + false alarms)—for the particular response group (e.g., Wixted et al., 2010; Ingram et al., 2011). Strictly speaking, the precise measure is defined as HR/(HR + FAR), where the HR (hit rate) is the ratio of hits in a category to the total number of targets, and the FAR (false alarm rate) is the ratio of FAs in a category to the total number of lures. Given that we used an equal number of targets and lures, computations using hits and FAs were identical to those using HR and FAR. Note that the frequency of items in different response groups is determined by participants' responses and is not experimentally controlled. Importantly, the measure of proportion correct does not vary as a function of this frequency.

In contrast, other potential measures (e.g., hits–FAs; HR–FAR; d′) fluctuate as a function of the frequency of items in the response group. To illustrate, if accuracy were estimated using (hits–FAs), then a response group which includes, say, twice as many items would yield a measure of accuracy twice as high (e.g., the accuracy for 30 hits and 4 FAs would be 26, as compared to an accuracy of 52, for 60 hits and 8 FAs). The measure of accuracy would likewise increase as a function of frequency of responses within the response group, had we used a measure of (HR–FAR).[2] To reiterate, the measure of proportion correct, i.e., hits/(hits + FAs), is not

---

[2] When comparing accuracy of overall performance across different conditions —as opposed to accuracy for a particular response group—measures of accuracy are often labeled 'sensitivity'. These include, but are not limited to, HR-FAR and d′. Unlike accuracy for a particular responses group, for overall performance, the frequency of items is fixed, in that it is solely determined by the experimenter and not dependent on participants' responses. These measures are thus the measures of choice to estimate overall sensitivity, but should not be used to estimate accuracy for a particular response group (see Rotello et al. (2015, 2008), for a discussion of considerations underlying the choice of particular measures of sensitivity).

affected by the frequency of items in the response group, which may fluctuate from participant to participant. Hence, it was the selected measure.

2.2.2. Exclusion analysis

To compute the zROC slope, we used the six-point confidence scale to generate five hypothetical response criteria for categorizing items as either ‘old’ or ‘new’. Thus, for the strictest criterion, only responses with the highest confidence (level 6) were considered ‘old’. Similarly, responses 5 and 6 were grouped together to represent a slightly more lenient hypothetical response criterion. At the other extreme was the grouping of responses 2–6, representing the most lenient response bias.

For each participant's five possible response criteria, hits and false alarms were transformed to scores in z-space. These scores were then plotted—for each of the five responses criteria—with hits as a function of false alarms, with linear regression yielding the zROC function. Across participants, the zROC function was linear, with mean $R^2=0.979$ ($SE=0.006$). This result does not comply with the DPSD model which predicts a U-shaped zROC function (Yonelinas and Parks, 2007). The mean slope value across participants was 0.72 ($SE=0.02$), and was significantly smaller than 1, $t_{15}=-12.43$, $p<0.001$. This result replicates similar slope values reported in the recognition literature (Diana et al., 2006; Glanzer et al., 1999; Rotello et al., 2004; Slotnick and Dodson, 2005; Wixted, 2007; Yonelinas, 2001).

Next we explored the changes in the zROC slope following exclusion of different response groups. Following Yonelinas (2001), we selectively excluded R responses in the three response groups of confidence levels 4, 5, and 6, which comprised 23.6% of the total trials (27% of the false alarms and 59% of the hits). The zROC slope increased after the exclusion of R responses (mean$=0.86$; $SE=0.04$), and the increase was significant, $t_{15}=-43.72$, $p<0.001$. Critically, however, in contrast to the results obtained by Yonelinas, the zROC slope across participants did not reach unity, and was significantly lower than 1, $t_{15}=-3.30$, $p=.005$. To reiterate, the unitary zROC slope has been cited as evidence in support of familiarity-based responses mediating performance, in both healthy individuals and amnesic patients. This is because according to the DPSD model, familiarity-based responses are represented by equal-variance signal detection distributions. However, the removal of R responses in our data did not increase the zROC slope to 1, thus failing to replicate the Yonelinas result.

Next, we explored the hypothesis that our failure to replicate the increase of the zROC slope to unity with the exclusion of R responses was because such increases necessitate the exclusion of strong memories—sparing only weak memories—rather than the exclusion of recollection-based responses—sparing familiarity-based responses. The strongest memories are defined as those corresponding to the response groups with the highest accuracy. To closely mimic the Yonelinas' exclusion scheme, we too excluded three response groups. Specifically, we excluded responses of the three most accurate response groups: 6R, 6K, and 5R (see Table 1).

Table 1
Accuracy collapsed across participants in Experiment 1. Interestingly, accuracy was associated with an ordering of response groups of R-K nested under confidence.

	4		5		6	
	Know	Remember	Know	Remember	Know	Remember
Accuracy	0.55	0.58	0.74	0.81	0.91	0.97

Interestingly, exclusion of the strongest response groups (23.3% of total trials; 14% of the false alarms and 61% of the hits)—resulted in an increase in the zROC slope to 0.95 ($SE=0.03$), $t_{15}=9.10$, $p<0.01$, which was not significantly different than 1, $t_{15}=-1.97$, $p=0.068$. Thus, the removal of strongest responses, rather than recollection-based responses, increased the zROC slope to 1.[3] Therefore, in our data, weak memories seem to be the ones which conform to equal-variance signal-detection distributions. Mean slopes are shown in Fig. 1.

In our final analysis, we examined whether an exclusion scheme of not-so-strong responses would still lead to an increase in the zROC slope, even to unity. Therefore, we excluded 5K category instead of the 5R category, yielding the exclusion of only a single R category—6R—and two K categories—6K and 5K. This new analyses yielded an increase in the zROC slope to 0.95 ($SE=0.04$), $t_{15}=6.39$, $p<0.01$, which was not significantly different than 1, $t_{15}=-1.38$, $p=0.189$.

2.2.3. Strength analysis

Throughout our analyses, we have suggested that the increase of the slope of the zROC curve was mediated by the fact that only weak memories remained following exclusion. It is a mathematical property of signal-detection, that removal of the distribution tail should yield target and lure means that are closer to each other (i.e., a decrease in d′). Still, we wished to show this directly, by comparing the zROC intercept before exclusion to that after the exclusion as well as between the exclusion schemes. The zROC intercept corresponds to d′, a measure of the distance between the means of target and lure distributions. As such, smaller distances (d's) represent weaker memories.

As expected, across participants, the zROC intercept before exclusion ($M=1.28$, $SE=0.09$) decreased to $M=0.88$ ($SE=0.08$) following the exclusion of the strongest items. To compare the intercepts, here and in all subsequent intercept comparisons, we used the sign test[4]—in effect, a binomial test with a null of $p=.5$. All participants showed the decrease following exclusion, making it significant when tested with a sign test, $p<0.001$. As expected, a similar result was obtained when removing the R responses, wherein the zROC intercept also decreased significantly ($M=0.94$, $SE=0.07$), with all participants showing the decrease, $p<0.001$. Thus, the classic increase in the zROC slope that was demonstrated by Yonelinas when excluding R responses can easily be interpreted as mediated by a UVSD mechanism, wherein the remaining responses represent weak traces.

Further examination of the zROC intercept scores provided additional evidence to the idea that the increases to 1 of the zROC slope are mediated by the weakness of the remaining memories. Thus, as expected, the removal of the strongest responses, which yielded an increase of the zROC slope to 1, resulted in a lower intercept scores ($M=0.88$, $SE=0.08$) than the removal of the R responses ($M=0.94$, $SE=0.07$), with 12 out of 16 participants showing the effect, $p=.028$. These results support a strength-based interpretation of the results obtained by Yonelinas.

---

[3] To reiterate, a strength interpretation only predicts an increase of the zROC slope, but does not grant a special status to a slope of 1. In fact, to obtain evidence in favor of a unitary zROC slope, a Bayes factor must be computed, showing evidence in favor of the null hypothesis of slope$=1$. Such an analysis was *not* reported by Yonelinas (2002) in his investigations of different exclusion schemes, though his DPSD model does grant a slope of 1 a special status. Neither was it reported in amnesic patients by Aggleton et al. (2005) and Yonelinas et al. (2002), though they too granted a slope of 1 a special status. Here, we did not compute a Bayes factor, in that it would misleadingly suggest a special status to a slope of 1, which a strength interpretation does not suggest. Also see Section 4.3.

(*footnote continued*)

[4] We could not use a t-test, because d′ scores do not conform to a normal distribution, for several reasons, among which, because they cannot attain negative values.

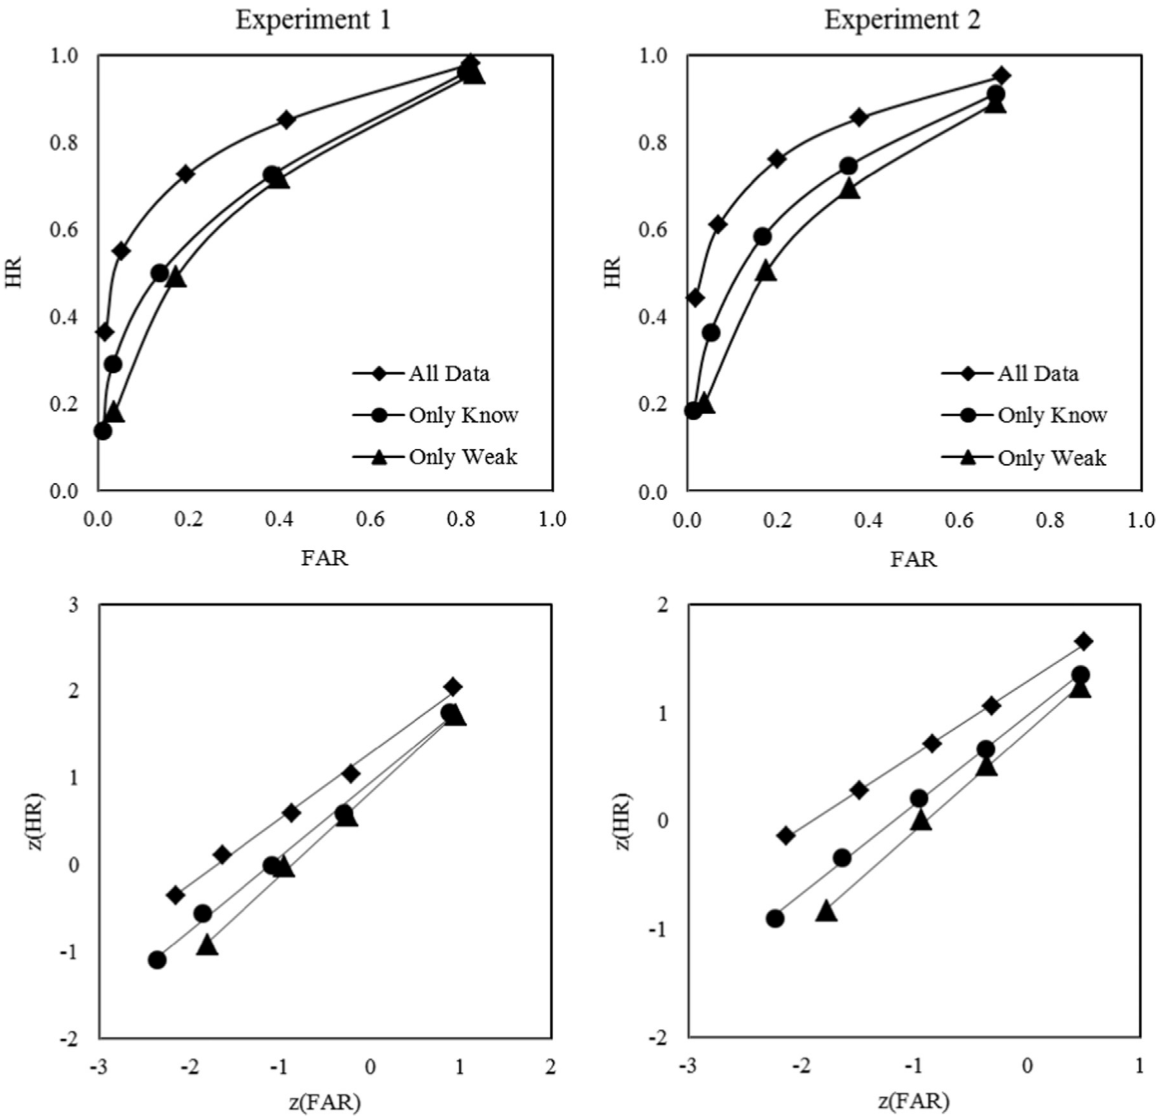

Fig. 1. Experiment 1 and 2 mean zROC slopes across participants. Points were averaged across participants.

Specifically, the increase in the zROC slope values seems to reflect the strength of the remaining memories, rather the lone presence of familiarity-based responses. Thus, when weaker memories remain, the target distribution more closely resembles that of the lures, in both mean (leading to a lower zROC intercept score) and variance (leading to a higher zROC slope).

Consistent with our findings, also in our final strength-based exclusion analysis (6R, 6K, 5K), which led to a zROC slope that was not significantly different than 1—the intercept of the zROC curve showed a marginally significant decrease after the exclusion ($M=0.83$, $SE=0.08$) in comparison with the exclusion of only R responses ($M=0.94$; SE=0.07), with 11 out of 16 participants showing the effect, $p=0.067$.

Taking together, our findings provide support for UVSD interpretation, according to which R and K do not map onto distinct processes, but may simply map onto strength. Accordingly, the unitary zROC slope reported in amnesic patients may, in fact, not be a signature of impairment to the recollective process but rather, a function of weak memory traces.

2.2.4. Auxiliary analyses: accuracy as a function of R-K nested within confidence

The current study was motivated by the idea that a strong association exists between R responses and strength (Dunn, 2004). This association was the motivating factor behind our endeavor to discover what mediated the increase of the zROC slope to 1, removal of R responses or removal of the strongest responses. Surprisingly, as it turned out, the association between R and strength was not as strong as we had assumed.

As can be seen in Table 1, an association emerged between accuracy and response groups, wherein accuracy decreased between the response groups (6R, 6K, 5R, 5K, 4R, 4K), oscillating between R and K at confidence levels 6, 5 and 4. We label this ordering of response groups as *R-K nested within confidence*. This linear decrease has never been reported in the recognition literature.

Importantly, this finding poses a challenge to the two classic recognition models. Specifically, UVSD suggests that R judgments reflect strong mnemonic strength, and hence should be coupled with high-confidence and high accuracy. This model, therefore, would find it difficult to defend a strong correlation as we observed between accuracy and R-K nested within confidence. Such a correlation would likewise compromise the DPSD model, which assumes that people recollected items reflect veridical memories, corresponding to both high confidence and high accuracy (but see Higam and Vockey, 2004; Roediger and McDermott, 1995; Wixted and Stretch, 2004).

To examine the extent of this serendipitous finding, we submitted the data to a simple linear-regression analysis on the collapsed data of all participants, examining accuracy as a function of the R-K nested within confidence ordering (i.e., we regressed a serial order number of six response groups (i.e., 1, 2, 3, 4, 5, 6) corresponding to the ascending order of accuracy we found: 4K, 4R, 5K, 5R, 6K, 6R). The coefficient of determination was $R^2=.977$, $F(1, 4)=177.24$, $p<0.001$. Additionally, we calculated the coefficient of determination for each participant, and obtained a mean $R^2=0.615$, which was significantly higher than 0, $t_{15}=13.13$, $p<0.001$.

We also examined the fit of the data to the UVSD and DPSD models. Both models predict that accuracy should be associated with *confidence nested within R-K* (i.e., 6R, 5R, 4R, 6K, 5K, 4K) ordering of response groups, that is, higher for all R responses than for all K response groups, with a decrease in accuracy associated with a decrease in confidence levels nested within R-K. An examination of the shared prediction of the two models to a confidence nested within R-K, we found a fit of only $R^2=0.384$, $F(1, 4)=2.49$, $p=.190$ for the aggregated data, and $R^2=0.362$ across participants. Inconsistent with both UVSD and DPSD, the linearity was found to be significantly higher for R-K nested within confidence than for confidence nested within R-K, $t_{15}=3.17$, $p=.006$. In Experiment 2, we explore a possible UVSD interpretation for this pattern in terms of idiosyncratic strategy that participants may have adopted in making their R-K judgments.

2.2.5. Spread of responses

To complete the picture, we examined the spread of subjective judgments across the different levels of confidence. Table 2 presents the proportions of R and K judgments for each of the confidence levels (4–6) categorized according to correctness of responses (hits, false alarms). Proportions were calculated out of the total'old' responses.

Examination of the table revealed two important patterns. First, both R and K responses spanned throughout the range of confidence levels 4–6. Note that while the highest level of confidence (level 6) drew the highest proportion of R responses across hits and FAs, a non-trivial percent of R responses—31%[5]—were associated with lower levels of confidence (levels 4–5). Additionally, 15% of K responses were associated with the highest level of confidence (level 6). Second, 27% of all false alarms were judged as R. Interestingly, both patterns contrast those of DPSD (Yonelinas, 2001), who found R responses to be made almost exclusively at the highest level of confidence, and who found only a negligible number of FAs that were judged R. We discuss this spread of R responses in the General Discussion, Section 4.1.

2.3. Discussion

In Experiment 1, we found that removal of R responses led to an increase of the zROC slope. However, despite the increase, the slope remained significantly below 1, thus failing to replicate Yonelinas (2001). According to Yonelinas, the exclusion of R responses from ROC analysis reflects the exclusion of recollection-based responses, leaving only familiarity-based responses in the analysis. As familiarity is assumed to adhere to equal-variance distributions according to the DPSD, removal of R responses should have increased the zROC slope to unitary—mimicking the unitary zROC slopes found in amnesic patients. Failure to obtain such a result in healthy participants suggests that deficit in mechanisms other than recollective processing, may be responsible for amnesia.

In contrast, an increase to unity of the zROC slope was observed, under exclusion schemes that did not reflect the removal of recollection-based responses. Specifically, when the response groups with the highest accuracy—not only R responses, but K responses as well—were excluded from the analysis, the zROC slope increased to unitary. Indeed, the idea that the exclusion of strong responses leads to an increase of the zROC slope is also supported by increase of the zROC slope under the exclusion of R responses (albeit not to unity). The increase we found likely reflects the exclusion of largely strong memories, although not the strongest. Together, our findings lend support to the notion that the increase of the zROC slope to 1 when R responses were excluded from the analysis may be due to removal of the strongest responses, which were associated, but not identical, to R responses (Dunn, 2004; Wixted and Stretch, 2004). If so, then a critical piece of evidence in support of the idea that amnesia reflects a deficit in recollective processes, as witnessed in the zROC slope of 1, may warrant further consideration (for a consideration of other evidence, see Section 4.5). At bottom, our investigation of the increase to unity of the zROC slope in healthy participants suggests that—in accordance with UVSD—in amnesia too, the unitary slope may simply reflect the existence of only weak memories. In Experiment 2, we further explored our findings.

Table 2
Distribution of Experiment-1 hits and false alarms, categorized according to confidence level (4–6) and judgment type (Remember, Know), computed from of the total number of 'old' responses (=3524). Data represent responses from 16 participants.

	Hits			False alarms		
	4	5	6	4	5	6
Remember	0.05	0.06	0.34	0.03	0.02	0.01
Know	0.14	0.11	0.07	0.12	0.04	0.01

3. Experiment 2

In Experiment 1, accuracy decreased as a function of R-K nested within confidence ordering of response groups (6R, 6K, 5R, 5K, 4R, 4K). This pattern poses a challenge for DPSD and UVSD, both which predict that accuracy should be correlated with confidence nested within R-K (that is, that all R responses should be more accurate that all K responses). We now wish to address the possibility that the structure of the task itself—confidence rating followed by R-K judgments—may have had an effect on the observed pattern of results.

We see no apparent recourse for DPSD to account for this finding (but see Section 4.2). In contrast, the decrease in accuracy as a function of response category can be accommodated by UVSD, if participants construed the *r*-K task as another strength judgment, 'zooming in' on their initial confidence ratings. Specifically, when initially rating their confidence, participants presumably set two criteria—corresponding to confidence levels 4, 5 and 6—along a hypothetical mnemonic-strength continuum, to distinguish between three levels of strength. Once they were in 'criteria-setting mode', it is feasible that participants construed the subsequent R-K task as that of 'zooming in' on finer grains of strength. Accordingly, they may have set additional fine-grained strength criteria, thereby creating subdivisions with the initial strength categories, for a total of six strength categories. Thus, according to our suggestion, it is possible that participants classified items as stronger (R) and the weaker (K) nested within each confidence level on an ad-hoc basis, yielding the observed pattern of accuracy correlated with R-K nested within confidence (6R, 6K, 5R, 5K, 4R, 4K). In Experiment 2, we tested this suggestion of an idiosyncratic 'zooming in' strategy. We also wished to replicate our results regarding the exclusion of strong- vs. R-based responses.

---

[5] Across hits and FAs, calculated as [Remember (4+5)/Remember (4+5+6)].

To this end, we reversed the order of the two tasks, with participants first making the remember-know-new judgments, and only subsequently, providing confidence judgments on a 3-point scale. This paradigm provided us with the same six response groups, as in Experiment 1, which were the focus of the present investigation (i.e., confidence 4–6 crossed with R-K). If the ad-hoc ‘zooming in’ interpretation of results is correct, then under the reverse ordering of tasks, participants should initially set the criteria to divide items to ones with the strongest memory traces (R responses), moderately-strong traces (K responses) and the weakest (‘new’ responses) memory traces. Subsequently, in the confidence-ratings task, these response groups would be refined based on the items' strength,[6] creating sub-division of three confidence levels. Thus, the strongest response groups—associated with the highest accuracy—should be all the R response groups, and these would be stronger than all the K response groups—associated with lower accuracy. Thus, Experiment 2's goal was to test this possible strength-based interpretation of the R-K task—this supporting UVSD—whereby, confidence would now be nested within R-K (R6, R5, R4, K6, K5, K4). If however, R-K would still emerge as nested within confidence, the data would pose would pose a strong constraint on any model of recognition, challenging at once both DPSD and UVSD.

3.1. Material and methods

3.1.1. Participants

Twenty one students (4 males, aged 22.4 ± 1.2 years old) participated in the experiment for credit in an undergraduate psychology course at Tel-Aviv University. One participant was excluded from the analysis because she did not have any FA in response groups 5 and 6, preventing us from calculating ROC curves. Two additional participants were excluded based on the Experiment-1 exclusion criteria, thus excluding participants who had more than 40% of their responses at a single confidence level as well as had less than 15% of the responses categorized as K judgments (thus, leading to two response groups, 6K and 5K, which comprised less than 0.03% of trials).

3.1.2. Stimuli and procedure

The apparatus, stimuli and procedure were identical to those of Experiment 1, with the following change: Task order was reversed. Hence, first a ‘Remember-know-New’ judgment was presented. Thereafter, for each item judged either ‘old’ or ‘new’, a subsequent screen was presented, in which participants made recognition judgments using a three-point confidence scale. Participants were presented with confidence levels 4–6 for R or K responses, with the identical labels as in Experiment 1. For ‘new’ responses, participants were presented with confidence levels 1–3, again, with the identical Experiment-1 labels.

3.2. Results

3.2.1. General

Across participants, mean performance yielded a 75% hit rate ($SE=2.2\%$) and a 21% false alarm rate ($SE=2.3\%$). More specifically, R judgments reflected a mean 40% HR and mean 3% FAR, whereas K judgments were associated with 35% HR and 18% FAR. Across hits and FAs, R judgments were associated with higher confidence ratings ($M=5.67$; $SE=0.04$) than K judgments ($M=4.81$; $SE=0.08$) $t_{17}=10.22$, $p<0.001$. In addition, in terms of recognition accuracy of studied items, the proportion correct for R judgments was higher ($M=92\%$; $SE=0.01$) than for K judgments ($M=67\%$; $SE=0.03$), $t_{17}=9.31$, $p<0.001$.

Table 3
Accuracy per response group collapsed across participants in Experiment 2. Like in Experiment 1, the results reveal that accuracy was associated with an ordering of response groups of R-K nested under confidence. This pattern challenge both DPSD and UVSD.

4		5		6	
Know	Remember	Know	Remember	Know	Remember
0.51	0.61	0.69	0.84	0.90	0.98

3.2.2. Accuracy analysis

We calculated the accuracy—proportion correct—for each of the six response groups created by the crossing judgment type (R, K) with confidence level (4–6). Table 3 presents the probability of providing a correct recognition judgment for each of the six groups. As in Experiment 1, accuracy was associated with an ordering of response groups of R-K nested under confidence. Importantly, because our idiosyncratic ‘zooming in’ interpretation could unlikely apply to this pattern, these data challenge not only DPSD but UVSD as well.

Linear-regression analysis on the aggregated data of all participants, as a function of the R-K nest within confidence (6R, 6K, 5R, 5K, 4R, 4K) yielded a coefficient of determination $R^2=0.987$, $F(1, 4)=374.2$, $p<0.001$. Additionally, we calculated the coefficient of determination for each participant and obtained a mean $R^2=0.628$, which was significantly higher than 0, $t_{17}=16.39$, $p<0.001$. Next, we reordered the response groups according to the UVSD and DPSD model predictions, with confidence nested within R-K response groups (i.e., 6R, 5R, 4R, 6K, 5K, and 4K). This produced a non-significant fit of only $R^2=0.481$, $F(1, 4)=3.91$, $p=.119$ for the aggregated data and of $R^2=0.495$ across participants. When comparing between models, linearity was found to be marginally higher for R-K nested within confidence, thus challenging both DPSD and UVSD ($t_{17}=1.97$, $p=.06$).

3.2.3. Exclusion analysis

Next we turned to the most critical analysis, arguably an experimentally-induced amnesia. As in Experiment 1, the zROC function was linear, with mean $R^2=0.921$ ($SE=0.014$), across participants. The mean slope value across participants was 0.70 ($SE=0.03$), which was significantly smaller than 1, ($t_{17}=9.616$, $p<0.001$).

Following Yonelinas (2001), we selectively excluded R responses, which comprised 21.5% of the total trials (16% of all false alarms and 53% of all the hits). The mean zROC slope across participants significantly increased to 0.85 ($SE=0.04$), $t_{17}=-32.01$, $p<0.001$. Critically, however, in replication of Experiment 1, the increase remained significantly different from unity, $t_{17}=9.616$, $p<0.001$. Thus, as in Experiment 1, here too we failed to replicate Yonelinas. Still, the increase of the zROC slope following the exclusion of R was consistent with the notion, that removing strong response, leads to an increase in the zROC slope.

Next, we explored the idea that our failure to replicate the increase of the zROC slope to unity with the exclusion of R responses was because such increases necessitate the exclusion of the strongest memories—including only weak memories in the analysis—rather than the exclusion of recollection-based responses—sparing familiarity-based responses. As in Experiment 1, we

---

[6] Note that the number of criteria is identical for Experiments 1 and 2. In Experiment 1, for the six-point confidence scale, five criteria are ostensibly set, and subsequently for the R-K task, an additional criterion for each of the three ‘old’ confidence levels, for a total of eight criteria. Likewise, in Experiment 2, a total of eight criteria would be set, two in the ‘remember-know-new’ task, and subsequently an additional six criteria—two criteria for each of the three possible responses in R-K-N task.

excluded three response groups corresponding to the response groups with the highest accuracy (6R, 6K, 5R; see Table 3), corresponding to the pattern of R-K nested within confidence. Similar to Experiment 1, exclusion of the responses (26.6% of the total trials; 15% of all false alarms and 66% of the hits) in the three strongest response groups (6R, 6K, 5R), resulted in an increase of the zROC slope to 0.97 ($SE=0.04$), $t_{17}=11.03$, $p<0.001$, which was not significantly different than 1, $t_{17}=0.68$, $p=.504$ (but see Footnote 3, where we reiterate that a strength interpretation does not grant a special status to a slope of 1). Mean zROC slopes are shown in Fig. 1.

Like Experiment 1, we undertook a more challenging examination, wherein we excluded the strongest responses, but instead of excluding the 5R category—which overlaps with the R exclusion scheme—we excluded the 5K category instead. Thus, we excluded only a single R category—6R—and two K categories—6K and 5K. Also this analysis yielded an increase in the zROC slope to 0.94 ($SE=0.04$), $t_{17}=5.89$, $p<0.001$, which was not significantly different from 1, $t_{17}=1.815$, $p=.087$.

3.2.3.1. Analysis across Experiments 1 and 2. Finally, looking at both experiments, no significant difference was found between the experiments. Specifically, no difference was found between Experiment 1 and Experiment 2 in the mean zROC slopes calculated for the entire data ($t_{32}=0.52$, $p=.609$), for the data under R exclusion scheme ($t_{32}=0.32$, $p=.753$), or for the data under the exclusion scheme of the most accurate responses ($t_{32}=0.56$, $p=0.578$). Additionally, combining the data from both experiments did not change the results found in the individual experiments. Specifically, the zROC slope for the joint analysis ($M=0.70$, $SE=0.02$) increased when R responses were removed from the analysis ($M=0.85$, $SE=0.03$), $t_{33}=6.58$, $p<0.001$. However, the increase was found to be significantly smaller than 1, $t_{33}=5.01$, $p<0.001$, thus failing to replicate Yonelinas. Similarly, the removal of the strongest response groups from the analysis (6R, 6K, 5R) increased the zROC slope to 0.96 ($SE=0.02$), $t_{33}=14.44$, $p<0.001$, which was not significantly different than unity, $t_{33}=1.74$, $p=0.092$.

3.2.4. Strength analysis

Though it is a mathematical property of signal-detection theory, we wished to directly demonstrate how the increase of the zROC slope was mediated by the fact that the remaining responses represented only the weaker memories, as in Experiment 1. As anticipated, under both exclusion schemes, the zROC intercept ($M=1.29$, $SE=0.08$) decreased, across participants, following the exclusion of responses. Specifically, the zROC intercept decreased when the strongest items were removed ($M=0.86$, $SE=0.07$), with all participants showing the effect, $p<0.001$. A similar result was obtained when removing the R responses ($M=0.99$, $SE=0.08$), with all participants showing the effect, $p<0.001$.

In addition, we predicted that exclusion scheme that yielded slopes (statistically) equal to 1—presumably, reflecting the removal of the very strong memories—would yield a significantly larger decreases to their d’, as compared to exclusion scheme that yielded slopes that remained smaller than 1. This prediction was borne out by our data. The increase of the zROC slope to 1 when the strongest responses were excluded, but not when R responses were excluded, was mirrored for the intercept scores. Thus, the intercept score showed a significant decrease when the strongest responses were removed ($M=0.86$, $SE=0.07$) than when R responses were excluded ($M=0.99$, $SE=0.08$), for 16 out of 18 participants, $p<0.001$. Similarly, after the exclusion of response groups (6R, 6K and 5K), we saw a greater decrease to the intercept of the zROC curve ($M=0.83$; $SE=0.06$) than after the exclusion or R responses ($M=0.99$, $SE=0.08$), with 15 out of 16 participants showing the effect $p=0.003$.

Taken together, when examining the correlation between the mean zROC slope and its mean intercept under the different exclusion schemes in both experiments,[7] a correlation of $r_6=-0.96$, $p<0.001$ was found. This reinforces our suggestion that the increase of the slope to 1 is mediated by the fact that remaining memory traces were weak (cf., Glanzer et al., 1999). These findings thus provide support for a strength-based interpretation of the increase in the zROC slope. Accordingly, an increase to the unitary zROC slope reported in amnesic patients is, in fact, not a signature of impairment to the recollective process. Rather, it may reflect the inability to boost the strength of any memories, neither recollection nor familiarity.

In summary, Yonelinas (2001) report of a zROC slope no different from 1 following the removal of R responses was not replicated. Moreover, an increase to 1 was found under exclusion schemes that were not directed exclusively at R judgments, but included K responses as well. We interpret the increase of the zROC slope to the notion that only the weaker memories remained. The idea of sparing of only weak memories is well reflected by significant decrease in the distance between the target and lure distributions, indexed by a decrease in the zROC intercept. In the General Discussion (Section 4.1) we argue that these finding justify a reevaluation of the notion that the functional deficit underlying amnesia is an impairment to recollective processes. In fact, we will argue that it reflects weak memories, recollection as well as familiarity.

3.2.5. Spread of responses

Like in Experiment 1, we examined the spread of subjective judgments across the different levels of confidence. Table 4 presents the proportions of R and K judgments for each of the confidence levels (4–6) categorized according to correctness of responses (hits, false alarms). Proportions were calculated out of the total ‘old’ responses.

Examination of the table revealed almost identical patterns to that of Experiment 1 (Table 2). This provides a possible indication that the reversal of task order did not affect performance in any meaningful way. Indeed, the two patterns observed in Experiment 1, were replicated here. First, both R and K responses spanned throughout the range of confidence levels 4–6. Note that while the highest level of confidence (level 6) drew the highest proportion of R responses across hits and FAs, a non-trivial percent of R responses—24%[8]—were associated with lower levels of confidence (levels 4–5). Additionally, 26% of K responses, hits and FAs, were associated with the highest level of confidence (level 6). Second, 17% of all false alarms were judged as R. As noted in Experiment 1, both patterns contrast those of Yonelinas, 2001, who found remember responses to be made almost exclusively at the highest level of confidence, and who found only a negligible number of FAs that were judged R. In the General Discussion, we discuss the interpretation of the observed spread of R responses found in Experiments 1 and 2.

4. General discussion

In this article, we reported two novel results regarding recognition memory. First, that accuracy—a proxy of memory strength—decreased with an ordering of response groups of R-K nested within confidence (6R, 6K, 5R, 5K, 4R, 4K). Second, that the exclusion of the most accurate memories yielded a zROC slope of 1,

---

[7] Eight groups were included, four from each experiment: (1) data without exclusion (2) removal of R responses (3) removal of groups 6R, 6K, 5R (4) removal of groups 6R, 6K, 5K.

[8] Across hits and FA, Remember (4+5)/Remember (4+5+6).

Table 4
Distribution of Experiment-2 hits and false alarms, categorized according to confidence level (4–6) and judgment type (Remember, Know), computed from of the total number of 'old' responses (=4171). Data represent responses from 18 participants.

	Hits			False alarms		
	4	5	6	4	5	6
Remember	0.03	0.05	0.33	0.02	0.01	0.01
Know	0.13	0.11	0.13	0.12	0.05	0.01

reflecting equal variance of the target and lure distributions. We revealed these results while putting to test the notion that the nature of the functional deficit in amnesia is in recollective processing (Gilboa et al., 2006; Hirst et al., 1988; Schacter et al., 1984; Turriziani et al., 2008; Yonelinas et al., 1998).

One key finding that has helped support the notion of a recollective deficit in amnesia is that these patients' zROC analysis yields a unitary zROC slope, reflecting an equal-variance target and lure distribution. This unitary zROC slope contrasts the smaller-than-one zROC slope found in healthy participants, reflecting an inequality of variance—shorthand for the larger variance of the target distribution than that of the lure distribution. DPSD (Yonelinas, 1994) proposes that the inequality of variance found in healthy participants is an outcome of a mixture of an equal-variance familiarity signal together with a binary recollection signal. DPSD argues that if familiarity reflects an equal variance SDT process, then patients with severe recollection impairments should produce unitary zROC slopes. That a unitary zROC slope was found in amnesia, has thus served to bolster the idea than amnesia is a deficit in recollective processing (for a discussion of further evidence, see Section 4.5).

Here, we challenged the zROC evidence for a recollective deficit in amnesia, by investigating the analogous unitary zROC slope found in healthy participants. Yonelinas (2001) reported that when R responses were excluded from ROC analysis—presumably reflecting the exclusion of recollection-based responses—unitary zROC slopes were found in healthy participants. Like the amnesic unitary zROC slope, also the one uncovered in healthy participants following the exclusion of R responses was thus cited by Yonelinas as evidence in support of DPSD and the postulated equal-variance familiarity signal.

In two experiments, the original Yonelinas finding was not replicated. The zROC slopes were immune to an increase to unity following the removal of R responses. Critically, in our data, when accuracy was used as a proxy to strength, R responses were not the strongest responses. Rather, the strongest memories comprised 6R, 6K and 5R (the first three responses groups of the R-K-nested-within-confidence ordering of response groups). It was only when we applied an exclusion scheme wherein these three responses groups were excluded, did an increase emerge of the zROC slope to unity.

Our findings thus open the possibility, previously endorsed by Squire et al. (2007), that the unitary zROC slope found in amnesia reflects the operation of only weak memory traces. As described in the introduction, converging evidence for the notion that zROC slope of 1 reflects weak memories was found in the decrease in the zROC slope which has been found to be associated with enhanced memory performance (Glanzer et al., 1999). This idea was also supported by the decrease in the zROC slopes of hippocampal-injured patients presented with short study lists, thereby enhancing their performance so as to match that of controls (Wais et al., 2006). Thus, our investigation of zROC slope under different exclusion schemes joins earlier findings in suggesting that the slope of the zROC slope, in healthy participants and possibly in amnesic patients, may better be characterized as reflecting recognition of weak memories, rather than as reflecting the involvement of recollective process in recognition decisions.

Our suggestion that a unitary zROC slope reflects the operation of weak memories is mute with regard to possible sources of such weak memories. Possible sources include, but are not limited to, frail encoding operations, poor retrieval operations, an inability to recapitulate encoding operations at retrieval, and impaired maintenance of the memories during storage. Our suggestion only addresses the final outcome—the memories of these patients are weak, thereby yielding unitary zROC slopes.

Critically, we argue that it is not sufficient to characterize performance—of amnesics or of healthy individuals—only in terms of weak versus strong memory traces. This is because a unidimensional strength signal (i.e., UVSD) cannot account for our finding of R-K nested within confidence (with the use of an idiosyncratic strategy ruled out in Experiment 2). It is noteworthy that our findings are not the first to challenge UVSD. For example, Rotello et al. (2004) found that two-point zROC slopes using RK data were very different from ROCs derived via old-new item recognition. If R judgments are assumed to be merely high-confidence old decisions, the same model must fit both the old–new data and the R-K data, which it did not. Consequently, the two-point ROC curves derived from RK data do not differ from item recognition just in confidence, i.e., strength. Rather, they must reflect some other form of evidence instead of, or in addition to, memory strength (see paragraphs below for additional evidence against UVSD). We now describe the CDP model, which we argue provides the best model to account for the plethora of RK recognition data, including those reported in this article.

4.1. *Evidence for the Continuous Dual-Process (CDP) model of recognition and its variants*

In Experiments 1 and 2, accuracy decreased linearly with R-K nested under confidence (6R, 6K, 5R, 5K, 4R, 4K). This nested pattern, which we had not predicted, is difficult to reconcile in terms of the classic UVSD model, according to which participants make R-K judgments by placing the criterion along the mnemonic-strength distribution, with the criterion placed in one, invariable position. Thus, within-participant, each of the R response groups is predicted to be stronger than each of the K response groups——an effect not observed in the data. Likewise, this very prediction—not borne out in the data—applies to the DPSD model, which assumes that all R responses represent direct access to veridical memories and should, therefore, be more accurate than all K responses. Below, in Section 4.2, we consider possible rebuttals from proponents of DPSD and UVSD to our findings.

Much to our surprise, the pattern of R-K nested under confidence has shown up in the literature before, though has gone completely unnoticed.[9] Specifically, it is evident in Figures 5, 8 and 9 of Ingram et al. (2011), it is evident in Figure 14 of Wixted and Mickes (2010), and it is largely evident in data reported by Rotello et al. (2005, Appendix A1), summed across the conservative and neutral conditions to reduce noise. This pattern therefore seems sufficiently systematic and reliable as to require an interpretation.

We suggest that the recent CDP model (Wixted and Mickes, 2010) can accommodate this nested pattern. Like DPSD, CDP proposes that recognition comprises both recollection and familiarity. However, CDP suggests that not only familiarity, but recollection too, is a continuous signal that conforms to a signal-detection mechanism, rather than the high-threshold recollection signal

---

[9] We thank John Wixted for pointing this out in a review of an earlier draft of this paper.

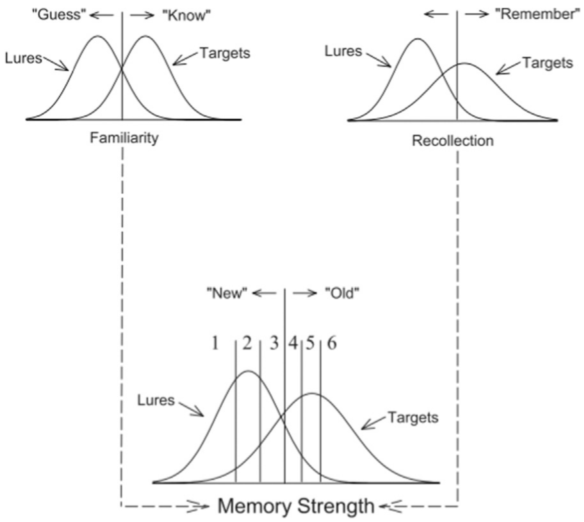

Fig. 2. An illustration of the continuous dual-process (CDP) model (taken from Wixted and Mickes (2010)). R-K judgments are based on the familiarity and recollection distributions whereas old/new judgments are based on the aggregated signal below.

proposed by DPSD. Specifically, like the UVSD model, CDP suggests that for most recognition decisions, a sum-of-evidence, mnemonic strength signal is consulted. However, CDP maintains that this signal is actually a composite of the two parent SDT signals—that of recollection and familiarity (see Fig. 2). Importantly, it is these 'parent' signals, both continuous in nature, which are accessed when making R-K judgments. Thus, according to CDP, both recollection and familiarity should be found associated with a wide range of confidence responses (e.g., levels 4–6), and can be found for target as well as lure items. This synthesis of the UVSD and DPSD models gives CDP much flexibility to handle constraining data patterns, including, as we shall see below, the nested pattern we observed Fig. 2.

Prior support for CDP comes from at least three findings, which can be accommodated by neither DPSD or UVSD. First, R and K responses have been found to be distributed across the various confidence levels (Rotello et al., 2005; Wixted and Mickes, 2010). Such a spread was also found in our data (see Tabless 2 and 4). Second, for many participants, the confidence levels associated with R and K responses overlaps by more than a single confidence level (Wixted and Stretch, 2004). This finding too is evident in our data.

Third, and most important to our present concerns, K responses have been found to be coupled with high confidence and highly accurate responses. Specifically, when examining accuracy at the highest confidence level (i.e., 20 on a scale of 1–20), K[10] responses were found to be more accurate than R responses at lower levels (16–19) (Ingram et al., 2011). Also this finding appeared in our data. Thus, although items were accurately recognized with high confidence, they were nevertheless not judged R, but rather, K. This finding is a laboratory analog to the well-known "butcher-in-the-bus" phenomenon (Mandler, 1980), wherein a strong sense of familiarity is sometimes elicited to a face, for which no episodic information can be recollected. The butcher-in-the-bus pattern poses a challenge not only to UVSD, but also for DPSD, which assumes that recollection-based responses should always be associated with higher confidence and higher accuracy than familiarity-based responses.

The existence of the butcher-in-the-bus pattern provides strong support to the notion that the individual memory signals of recollection and familiarity signals are continuous. Indeed, Wixted and Mickes (2010) ran a simulation study of their model, demonstrating that CDP can yield high confidence (i.e., confidence level 6, in a scale 1–6) K responses which are more accurate than R responses, made with lower confidence (i.e., level 5).[11]

Our finding that accuracy is associated with R-K nested under confidence levels 4–6 provides a replication of Ingram et al. (2011). Moreover, it extends their findings by demonstrating the complimentary pattern—when recollection acquires the weakest value. Thus, we found that recollection-based responses (4R) can be less accurate than familiarity-based responses (5K) even for somewhat weaker memories. Thus, our nested pattern established that K judgments with higher accuracy than R judgments, can be generalized to include not only K judgments made at the very highest level of confidence, but also weaker 5K responses (level 5), that though weaker, may still be more accurate than 'frailer'—level 4R responses.

At an intuitive level, the association between accuracy and the nested pattern can be obtained because R-K judgments and confidence judgments are based on different aspects of the same evidence. Thus, R-K judgments are based on the separate distributions of familiarity and recollection. In contrast, confidence judgments are based on the aggregate signal, wherein the strength of R and K can offset each other to various degrees, thereby producing a range of confidence levels. Because the level of confidence will always be a function of the original strength of the parent signals, high confidence judgments will typically correspond to stronger memories than lower confidence judgments, thus producing more accurate K responses (at level 6 or level 5) than R responses (at levels 5 and 4, respectively). Moreover, because the recollection distribution is on average, stronger than the familiarity distribution,[12] the model predicts that—within each confidence level—if the two parent distributions are consulted, R responses would on average be stronger, and thus more accurate, than K responses.

To establish that our observed pattern of R-K nested within confidence is indeed viable by CDP, we ran a set of CDP simulations, in which we either changed the Wixted and Mickes (2010) placement of the R criterion[13] or else changed the placement of one of two confidence criteria.[14] Our simulations revealed that CDP can yield a pattern wherein accuracy is most strongly associated with a R-K nested under confidence. Indeed, of the existing dual-process models in the recognition literature which address R-K performance, the CDP model seems to be the only candidate that is able to interpret a robust correlation between accuracy and an interlaced ordering of response groups.[15] Still, it is noteworthy that while our results are possible by CDP, the nested pattern is not the only pattern possible under CDP. Future research is required to delineate why this pattern turned up in our experiments as well is

---

[10] These authors used a remember-familiar task, which is identical in everything but the verbal label given to non-remembered items.

[11] To obtain the butcher-in-the-bus pattern, Wixted and Mickes (2010) placed the R criterion 1.8 SDs above the mean of the recollection lure distribution and placed the K criterion 1.5 SDs above the mean of the familiarity lure distribution. The five confidence criteria were placed in relation to the mean of the lure distribution in the following locations, 2.0, 1.5, 1.0, 0.50, and 0, with the units equal to the SD of the lure distribution ($\sim$1.41). Thus, in absolute values, these placements correspond to 2.82, 2.115, 1.410, 0.705, and 0, respectively.

[12] Although CDP can, in principle, accommodate the idea that the mean of the recollection distribution is lower than that of familiarity, the reverse order seems to be the more characteristic one, as implied by the values chosen in the CDP simulations in Wixted and Mickes (2010).

[13] Only the placement of the R criterion was changed. We changed it from 1.8 standard deviations above the mean of the recollection lure distribution (see Footnote 7, above), to 1.5 SDs.

[14] In relation to the mean of the lure distribution described in Footnote 13 (in units equal to the sd of the lure distribution), we either only changed the criterion placed at 1–0.709, or only changed the criterion placed at 1.5–1.205 (in absolute values, we changes to 1 and 1.7, respectively).

[15] Thus, implementing the changes described in either Footnote 13 or Footnote 14, yielded not only more accurate 5R than 6K, but also more accurate 4R than 5K.

in the studies cited above, and what are the necessary and sufficient conditions for obtaining this pattern.

The CDP-model interpretation was further supported by our exclusion analysis, where removal of the strongest responses, comprising both R and K, was found to be a more reliable procedure to produce equal-variance distributions than removal of R responses alone. This exclusion can be theoretically justified by the CDP notion that both the familiarity and the recollection signals can be strong. Taken together, it seems that equal-variance distributions more likely represent the distributions of weak memories than the distributions of familiarity-based responses. It is the removal of both strong signals—familiarity as well as recollection—that yielded equal-variance distributions, which are one of the signs, together with lower d', of weak memories. The implications of these conclusions to amnesia are discussed in Section 4.5.

We have this far interpreted our data in terms of the CDP model. However, based on recent neuroscientific evidence, several variants of this model were articulated. For example, the contextual-information account (Rugg & Vilberg, 2013: Rugg et al., 2012) suggests that hippocampal activity during retrieval reflects the amount of contextual information that is associated with the test item, rather than the subjective sense of recollection or the strength of an undifferentiated memory signal. On the other hand, others maintain that the hippocampus might relate to both familiarity and recollection information (Smith et al., 2011; Song et al., 2011). Additionally, it has recently been suggested that reinstatement of the neurocognitive processes engaged when an episode was encoded is critical for retrieval, and that this reinstatement is evident during both familiarity-based and recollection-based judgments, with recollection reflecting a continuous neural signal (Johnson et al., 2009; Leiker and Johnson, 2014). Thus, these models too view the recollective signal as continuous and can easily interpret our pattern of R-K nested within confidence.

4.2. Failure to replicate Yonelinas pattern: task instructions?

In Experiments 1 and 2, we failed to replicate two related findings. First, we found an interlaced ordering of responses (see Section 4.1), thereby failing to replicate the Yonelinas pattern, wherein R responses were stronger—more accurate—than K response. The notion that R responses should be stronger than K responses is predicted by both DPSD and UVSD (Wixted et al., 2004). In assuming that Rs are stronger than Ks, both DPSD and UVSD anticipate an increase to 1 of the zROC slope following R exclusions, interpreting this exclusion to reflect, respectively, removal of recollection and of the strongest responses. Second, we failed to replicate the increase to 1 of the zROC slope when R responses were excluded (Yonelinas, 2001). Our failure to replicate the increase to 1 following the exclusion of R[16] is obviously a function of the fact that in our experiments, the strongest responses comprised a blend of R and K responses. It was the removal of this blend of responses, not that of Rs alone, which led to a unitary zROC slope.

How would DPSD and UVSD account for our findings, wherein the strongest responses included a blend of Rs and Ks? With regard to DPSD, one critical difference between Yonelinas (2001) and the current study are the instructions for the R-K task. The instructions used by Yonelinas required participants to respond R only if they could actually describe to the experimenter specific details about the experience of studying the word. In contrast, in our experiments, participants were informed that the two types of judgments (confidence and R-K) were unrelated. No additional information was provided and no suggestion was made as to the relative distributions of the two types of judgments. We suggest that the Yonelinas instructions may have led participants to mimic a hypothetical high-threshold mechanism, akin to that postulated by DPSD to reflect recollection, wherein an R response is given whenever any recollective detail is retrieved. Thus, these instructions may have steered participants away from the default mode of recollective processing, which we argue is not high-threshold.

To reiterate, participants' interpretation of the R judgments under the Yonelinas instructions may have been biased to reflect a high-threshold mechanism, though this is not the default mode of processing for R responses. It is possible, therefore, that under non-biasing instruction, R responses do not conform to a high-threshold process, and may better be characterized as a continuous process, as proposed by CDP. Still, as a rebuttal, Yonelinas may argue that it was our instructions, not his, that steered participants away from default processing of R responses. Thus, Yonelinas could argue that in our study, participants did not properly understand the RK instructions, and hence, their responses did not reflect the true nature of the recollection process.

That changes in instructions can lead to changes in response patterns was well demonstrated by Rotello et al. (2005). In their study, Rotello et al. compared the standard instructions used by Yonelinas to instructions wherein participants were explicitly informed that although remembering and knowing are different feelings of recognition, they both may vary in confidence (Rotello et al., 2004). The two versions of instructions produced different spread of responses and different ROC curves. Interestingly, the two versions of instructions correspond to specific theories. Thus, the Yonelinas instructions suggest that R responses reflect only strong, veridical memories (i.e., the DPSD model). In contrast, the Rotello et al. instructions suggest that both R and K responses are spread throughout the confidence scale, as predicted by their model (i.e., the STREAK model). It is not surprising, therefore, that the changes in instructions between Yonelinas and Rotello et al., affected participants' response strategies, and consequently, the ensuing spread of responses.

Note, that in our experiments, participants were informed that the two types of judgments (confidence and R-K) were unrelated. By only stating that the judgments were unrelated, our goal was to obtain a (relatively) theory-neutral data set. To the extent that we succeeded, our data may perhaps be less biased by a specific strategy which reflects a specific recognition model.

How would UVSD account for our finding that strongest responses included a blend of Rs and Ks, rather than only R responses? In the discussion to Experiment 1, we suggested the decrease in accuracy as a function of response category can be accommodated by UVSD, if participants construed the *r*-K task as another judgment of strength, 'zooming in' on their initial confidence ratings to make fine-grained judgments of strength. To rule out this interpretation, in Experiment 2, participants first made an RK judgment, and only subsequently a confidence judgment. RK response groups were still nested under confidence. This finding makes it unlikely that participants zoomed in on their confidence ratings, in that confidence ratings were made only *after* RK judgments. Still, UVSD may suggest that although the task order lent confidence to be nested within RK, in their mind, the mental representation of strength was identical to that of

---

[16] In our data, incorrect Rs were not uncommon (27% and 16% of all FA in Experiment 1 and 2, respectively). Therefore, they were excluded from the analysis. Because incorrect Rs were rare in his data set, Yonelinas only excluded correct Rs. In Experiment 1, when only correct R responses were excluded, an increase to 1 of the zROC slope was found. Still, these correct Rs strongly overlapped with the strongest responses, making equivocal the interpretation of the increase (in Experiment 1, 83% of the correctly recognized responses at confidence level 6 were judged R. In Experiment 2, this number reached only 70%). Either way, in Experiment 2, even the exclusion of only correct Rs, did not yield an increase to 1 and the Yonelinas finding did.

Experiment 1, with RK nested within confidence. We see no reason why despite task order, participants would nonetheless represent strength such that RK judgements are construed to be fine-grained levels strength within confidence. Still, this interpretation cannot be entirely ruled out.

4.3. *How probable an outcome is the increase to unity of the zROC slope?*

If we assume, as we have done in this article, that accuracy is an index of memory strength, then removing the most accurate responses is tantamount to the removal of the extreme responses from the target and lure distributions. According to the CDP model, participants may access the two distributions of recollection and of familiarity. Thus, extreme responses are removed from each of these distributions, according to the specific exclusion scheme (e.g., only R responses from the recollection distribution, or both R and K responses from both the recollection and familiarity distributions). This is consistent with signal-detection distributions, wherein the removal of extreme responses should lead to a reduction in the variability of target distributions and thus to an increase in the ratio between the variance of the lure and the target distributions, even up to unity.

Because an increase in the zROC slope is expected under removal of extreme observations, it seems remarkable that rather than simply observing an increase in our data, we found an increase to slope which was not statistically different from 1 (but see Footnote 3). We thus asked how probable it was to find an increase to unity when removing the most extreme responses for ROC analysis. To examine whether, when removing the most extreme responses, the increase to unity is a probable outcome, we ran several Monte Carlo simulations. For simplicity sake, we simulated an unequal variance signal-detection model with three parameter sets.[17] The first set included (values used in Wixted et al., 2004; Appendix B) and the remaining two sets were extracted from Experiments 1 and 2, based on the d' and the proportion of false alarms in each confidence level, averaged across participants. The parameters comprised the SD of the target distribution (assuming a SD of 1 for the lure distribution), the distance between the mean of the target distribution from that of the lure distribution in units of SD of lure distribution, and the location of the five confidence criteria, in units of the SD of the lure distribution (see Table 5 for the parameter values).

For each parameter set, we generated 10,000 distributions. From each distribution, we excluded one of eight predefined percentages of the observations (5%, 10%, 15%, …, 40%). Specifically, for 1250 distributions, the zROC slope was calculated after exclusion of the most extreme 5% of responses (both hits and false alarms), for a different 1250 distributions, the slope was calculated after the exclusion of the most extreme 10% of responses, and so on, in steps of 5%, until, for the last 1250 distributions, the zROC slope was calculated for the exclusion of 40% of the data.

We thus obtained eight mean zROC slopes, for different levels of exclusion as can be seen in Fig. 3 below. For all three sets of parameters, the results of a zROC slope of 1 were highly probable. More precisely, in the parameter set used by Wixted and Stretch (2004), the slope first increased to 1 (SD=0.07) after removing as few as 15% of the data. The slope increased gradually in the next exclusion steps until reaching 1.2 (SD=0.2) when 40% of the data was removed. In the parameter sets extracted in both Experiment 1 and Experiment 2, the slope increased monotonically with every exclusion step, until reaching 1 (Experiment 1: SD=0.06; Experiment 2: SD=0.04) after removing 20% of the data. In the next exclusion steps, the changes in the slope were not monotonic and varied between 1.01 (SD=0.05) and 1.06 (SD=0.07) in Experiment 1 and between 1.0 (SD=0.03) and 1.11 (SD=0.05) in Experiment 2.

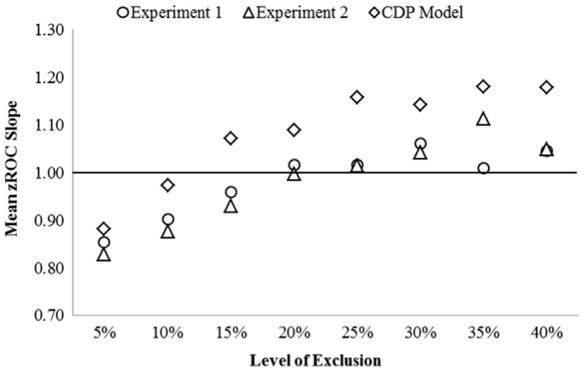

Fig. 3. Simulations of mean zROC slope at different levels of exclusions.

Table 5
Parameters set used in the simulations.

	Experiment 1	Experiment 2	Wixted and Stretch (2004)
SD_{Lure}	1	1	1
SD_{Target}	1.23	1.26	1.25
M_{Lure}	0	0	0
M_{Target}	1.83	2.02	1.65
Criterion 1	−.92	−.50	.93
Criterion 2	.21	.31	.98
Criterion 3	.87	.85	1.10
Criterion 4	1.65	1.49	1.60
Criterion 5	2.12	2.09	2.10

To reiterate, our exclusion of the strongest response groups in Experiment 1 and 2, reflected the removal of roughly 22–26% of the data. These simulations illustrate that the increase in the zROC slope in our data, to values which were not significantly smaller than 1, is not a surprising outcome after all.

4.4. *zROC analysis under deviations from normality*

It is a mathematical fact that when the underlying lure and target distributions are normal, then the slope of the zROC curve is a valid index to probe the ratio between variances. However, we used this slope to examine the ratio of the variances after the exclusion of the extreme responses. Exclusion of extreme responses, however, surely did not leave the remaining distribution as normal. So, how reliable is the measure of the zROC slope for distributions such as ours, where values at one tail of the distributions are systematically deleted?

To test the validity of the zROC slope in estimating the ratio between the standard deviation of the target and lure distributions when the underlying distributions are not normal (such as when responses are excluded from the extreme end of the distributions) we performed a simulation analysis. To this end, we calculated for each of generated distributions in each of eight exclusion steps (5%, 10%, 15%,…, 40%), the zROC slope as well as the ratio between the standard deviation of the generated target and lure distributions.

---

[17] All theories were formulated to account for the empirical finding of unequal variance. For this reason, we simulated an unequal variance signal detection distributions to examine whether given such a distribution, removal of the extreme distribution would yield a zROC slope of 1. Because according to the CDP model both the recollection and the familiarity distributions may conform to unequal-variance distributions, the results of such a simulation can be generalized to these distributions as well.

Indeed, for the parameter set of Experiment 1, the correlation between the average slope and standard-deviation ratio in the eight exclusion steps was found to be $r=0.95$ (and $r=0.83$ for the entire data set). Likewise, for the parameter set of Experiment 2, the correlation was found to be $r=0.97$ (and $r=0.90$ for the entire data set). Finally, for the parameter set of Wixted and Stretch (2004), the correlation was found to be $r=0.99$ (and $r=0.60$ for the entire data set). Thus, the zROC slope turns out to be a reliable estimate of the standard deviation ratio—even when the underlying distributions are noticeably not normal.

Another interesting point regarding the increase of the zROC slope is the parallel reduction in d' of the remaining data. In both Experiments 1 and 2, the increase in the zROC slope following the removal of responses, was correlated with a decrease in the d' of the remaining responses.[18] This suggests that the change in the zROC slope following removal of R responses reflects the distributions of the remaining weaker memories. Results supporting a high correlation between zROC slope and d' have previously been described (Glanzer et al., 1999). In their study, memory strength, probed by accuracy, was manipulated. Their study demonstrated that the zROC slopes were higher for less accurate conditions. Still, even in their least accurate condition, a slope as high as 1 was never observed. This makes our findings, whereby a slope of 1 is found, to be an important extension to those of Glanzer et al.

4.5. *The functional impairment underlying amnesia: Recollection or strong memories?*

In the introduction, we argued that support from zROC analysis that amnesia is a deficit in recollective processing entails the logical fallacy of inferring the consequent. Specifically, conditional on finding a unitary slope patients (or in K responses of healthy participants), evidence for a recollective deficit cannot be inferred. Indeed, our investigations of the zROC slope demonstrate how the existence of weak memories provides a more coherent interpretation of the increase found in the zROC slope.

To be sure, we presented no novel amnesic data in this article. Nevertheless, we suggest that the same interpretation applies to the unitary slope observed in amnesic patients. This suggestion is based on the fact that data for amnesics and the simulated data for healthy individuals seem to be identical. Both in amnesics and in healthy individuals, an increase in the zROC slope is observed to values which are not significantly different from 1 (for further details, see Footnote 3). We can offer no good reason to assume that a different mechanism mediates the identical increase of the zROC slope. Moreover, as we have stressed, the original justification of interpreting the unitary slope in amnesics was grounded on a logical fallacy so there is no need to embrace it in the face of new analyses.

Together, our investigations of the zROC slope and our observation of the interlaced pattern of responses suggest that functional deficit underlying amnesia may best be characterized as the inability to access strong memories. Amnesics are thus perhaps only able to access weak familiarity and recollection signals. This suggestion can be conceptualized in two different ways, describing different underlying mechanisms. First, that amnesia may spare weak memory signals (familiarity as well as recollection) while hurting strong traces. According to this idea, amnesia impairs an existing body of episodic information encoded in the brain. The second, that in amnesia, patients lose the neural mechanisms that enable the strengthening of the memory signals. According to this idea, amnesia is characterized by an inability to add new episodic information that can be likely retrieved, in that new memory traces, though formed, are very weak. Because anterograde amnesia—the focus of our investigations—is most typically characterized by an inability to learn new information, we think the second conceptualization may provide a more accurate portrayal of the mechanism mediating the amnesic impairment.

The original notion that amnesia can be characterized as an impairment in recollective processes was an outcome of not only the unitary zROC slope, but also of additional converging evidence. A full description of the evidence cited in favor of this notion is well beyond the scope of this article. Still, we would like to draw attention to one family of seemingly compelling findings, which is often cited as support for the notion of a recollective deficit.

Yonelinas et al. (1998) examined three experimental paradigms, and within each paradigm, obtained mathematical estimates for the values of familiarity and recollection. The three paradigms were the process-dissociation procedure (e.g., Bergerbest and Goshen-Gottstein, 2002; Jacoby, 1991; Verfaellie and Treadwell, 1993; Verfaellie, 1994; Yonelinas and Jacoby, 1995), the Remember-Know task (Gardiner, 1988; Knowlton and Squire, 1995; Schacter et al., 1997, 1996; Tulving, 1985, cf., Rosenstreich and Goshen-Gottstein, 2015) and ROC curves. Using all three paradigms, Yonelinas et al. obtained mathematical estimates using the DPSD model. Across all three paradigms, a substantial deficit in recollection was found, though accompanied by a small deficit in familiarity, when comparing performance to that of healthy individuals. In retrospect, the small deficit in familiarity should perhaps have provided a warning signal to the 'recollection-deficit' hypothesis.

Examination of the estimates of recollection for the amnesic patients revealed a reliable deficit in amnesia. Critically, however, the mathematical procedures used to derive the estimates all assumed stochastic independence between recollection and familiarity (for a description of the difference between functional and stochastic independence, see Moran et al., 2015). This assumption entails that given a high value of recollection, a high value of familiarity cannot be predicted. However, as stated by Norman and O'Reilly (2013, p. 612): 'there is no way to test this assumption using behavioral data alone because of chicken-and-egg problems (i.e., one needs to measure familiarity to assess its independence from recollection, but one needs to assume independence to measure familiarity).

Moreover, Bodner and colleagues (Brown and Bodner, 2011; Tousignant et al., 2012) asked participants to rate their experiences of both recollection and familiarity on different scales, using an independent rating task (Higham and Vokey, 2004). These authors found a positive correlation between recollection and familiarity. In addition, they found that some of the recollection-familiarity dissociations, which have been reported in the literature, disappeared when an independent rating task was used instead.

Finally, and most importantly, a recent study performed in our lab used the CDP model to estimate the magnitude of independence between recollection and familiarity (Moran and Goshen-Gottstein, 2015). It turns out, that these two signals were highly correlated (~ 0.6). Critically, with detailed examples, we showed how erroneously assuming independence introduces severe biases into the estimation procedure (for details, see Moran and Goshen-Gottstein, 2015). Taken together, it seems that parameter-estimation procedures which assume stochastic independence have all yielded biased results regarding the magnitude of parameters. Therefore, the conclusions of these procedures which state that amnesia is a deficit in recollective processing—

---

[18] Yonelinas and Parks (2007) suggest that the incorporation in the UVSD model of both a sensitivity parameter (d′) and a variability parameter (zROC slope) of the target distribution, implies that these parameters could be experimentally separated. However, weak memories, by definition, are hard to be distinguished from new items. Therefore, as we see in our data, weak memories have both similar variability and similar average strength to new items.

much like the finding of a unitary zROC slope in amnesics—must be viewed with great caution. In summary, if recollection and familiarity are indeed correlated, then a deficit in recollection should invariably be observed alongside a comparable deficit in familiarity—which is exactly the conclusion of the present investigation. Thus, all roads seem to lead to the same conclusion: Anterograde amnesia may reflect an inability to boost memories, both recollection and familiarity.

Acknowledgements

We thank John Wixted, Laura Mickes, Rani Moran and Yoav Kessler, for helpful discussions and thoughtful comments on an earlier version of this manuscript. We also thank Andrei Teodorescu for his help with the simulation work of Sections 4.3 and 4.4.

References

Aggleton, J.P., Vann, S.D., Denby, C., Dix, S., Mayes, A.R., Roberts, N., Yonelinas, A.P., 2005. Sparing of the familiarity component of recognition memory in a patient with hippocampal pathology. Neuropsychologia 43 (12), 1810–1823.. http://dx.doi.org/10.1016/j.neuropsychologia.2005.01.019.

Bergerbest, D., Goshen-Gottstein, Y., 2002. The origins of levels-of-processing effects in a conceptual test: evidence for automatic influences of memory from the process-dissociation procedure. Mem. Cognit. 30 (8), 1252–1262.

Bolognini, N., Ro, T., 2010. Transcranial magnetic stimulation: disrupting neural activity to alter and assess brain function. J. Neurosci.: Off. J. Soc. Neurosci. 30 (29), 9647–9650. http://dx.doi.org/10.1523/JNEUROSCI.1990-10.2010.

Brown, A.A., Bodner, G.E., 2011. Re-examining dissociations between remembering and knowing: binary judgments vs. independent ratings. J. Mem. Lang. 65 (2), 98–108. http://dx.doi.org/10.1016/j.jml.2011.04.003.

Cohen, N.J., Eichenbaum, H., 1993. Memory, Amnesia, and the Hippocampal System. MIT Press, Cambridge, MA.

Cohen, N.J., Poldrack, R.A., Eichenbaum, H., 1997. Memory for items and memory for relations in the procedural/declarative memory framework. Memory 5, 131–178.

Cohen, N., Squire, L.R., 1980. Preserved learning and retention of pattern analyzing skill in amnesia: dissociation of knowing how and knowing that. Science 210, 207–210.

Curran, T., 2000. Brain potentials of recollection and familiarity. Mem. Cognit. 28 (6), 923–938 http://www.ncbi.nlm.nih.gov/pubmed/11105518.

Davelaar, E.J., Goshen-Gottstein, Y., Ashkenazi, A., Haarmann, H.J., Usher, M., 2005. The demise of short-term memory revisited: empirical and computational investigations of recency effects. Psychol. Rev. 112 (1), 3–42. http://dx.doi.org/10.1037/0033-295X.112.1.3.

Diana, R.A., Reder, L.M., Arndt, J., Park, H., Arndtt, J., Park, H., Arndt, J., 2006. Models of recognition: a review of arguments in favor of a dual-process account. Psychol. Bull. Rev. 13 (1), 1–21. http://dx.doi.org/10.1016/j.bbi.2008.05.s010.

Donaldson, W., 1996. The role of decision processes in remembering and knowing. Mem. Cognit. 24 (4), 523–533 http://www.springerlink.com/index/H8569814461J6183.pdf.

Dunbar, K., Sussman, D., 1995. Toward a cognitive account of frontal lobe function: simulating frontal lobe deficits in normal subjects. Ann. N. Y. Acad. Sci. 769, 289–304.

Dunn, J.C., 2004. Remember-know: a matter of confidence. Psychol. Rev. 111 (2), 524–542. http://dx.doi.org/10.1037/0033-295X.111.2.524.

Dunn, J.C., 2008. The dimensionality of the remember-know task: a state-trace analysis. Psychol. Rev. 115 (2), 426–446. http://dx.doi.org/10.1037/0033-295X.115.2.426.

Eichenbaum, H., Yonelinas, A.P., Ranganath, C., 2007. The medial temporal lobe and recognition memory. Annu. Rev. Neurosci. 30, 123 http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2064941/.

Fortin, N.J., Wright, S.P., Eichenbaum, H., 2004. Recollection-like memory retrieval in rats is dependent on the hippocampus. Nature 431 (7005), 188–191. http://dx.doi.org/10.1038/nature02853.

Gardiner, J.M., 1988. Functional aspects of recollective experience. Mem. Cognit. 16 (4), 309–313 http://www.ncbi.nlm.nih.gov/pubmed/8750414.

Gilboa, A., Winocur, G., Rosenbaum, R.S., Poreh, A., Gao, F., Black, S.E., Moscovitch, M., 2006. Hippocampal contributions to recollection in retrograde and anterograde amnesia. Hippocampus 16, 966–980.

Glanzer, M., Kim, K., Hilford, A., Adams, J.K., 1999. Slope of the receiver-operating characteristic in recognition memory. J. Exp. Psychol.: Learn. Mem. Cognit. 25 (2), 500–513. http://dx.doi.org/10.1037//0278-7393.25.2.500.

Goshen-Gottstein, Y., Moscovitch, M., Melo, B., 2000. Intact implicit memory for newly formed verbal associations in amnesic patients following single study trials. Neuropsychology 14 (4), 570–578 http://psycnet.apa.org/journals/neu/14/4/570.

Hardt, O., Wang, S.-H., Nader, K., 2009. Storage or retrieval deficit: the yin and yang of amnesia. Learn. Mem. 16, 224–230.

Higham, P.A., Vokey, J.R., 2004. Illusory recollection and dual-process models of recognition memory. Q. J. Exp. Psychol. 57 (4), 714–744. http://dx.doi.org/10.1080/02724980343000468.

Hirst, W., Johnson, M.K., Phelps, E.A., Volpe, B.T., 1988. More on recognition and recall in amnesics. J. Exp. Psychol.: Learn. Mem. Cognit. 14 (4), 758–762.

Ingram, K.M., Mickes, L., Wixted, J.T., 2011. Recollection can be weak and familiarity can be strong. J. Exp. Psychol.: Learn. Mem. Cognit. . http://dx.doi.org/10.1037/a0025483

Jacoby, L.L., 1991. A process dissociation framework: separating automatic from intentional uses of memory. J. Mem. Lang. 30 (5), 513–541. http://dx.doi.org/10.1016/0749-596X(91)90025-F.

Johnson, J.D., McDuff, S.G.R., Rugg, M.D., Norman, K.A., 2009. Recollection, familiarity, and cortical reinstatement: a multivoxel pattern analysis. Neuron 63 (5), 697–708. http://dx.doi.org/10.1016/j.neuron.2009.08.011.

Knowlton, B.J., Squire, L.R., 1995. Remembering and knowing: two different expressions of declarative memory. J. Exp. Psychol.: Learn., Mem. Cognit. 21 (3), 699–710.

Leiker, E.K., Johnson, J.D., 2014. Neural reinstatement and the amount of information recollected. Brain Res. 1582, 125–138. http://dx.doi.org/10.1016/j.brainres.2014.07.026.

Macmillan, N.A., Creelman, C.D., 1991. Detection Theory: A User's Guide. Lawrence Erlbaum, New York.

Mandler, G., 1980. Recognizing: the judgment of previous occurrence. Psychol. Rev. 87 (3), 252–271 http://psycnet.apa.org/journals/rev/87/3/252.

Manns, J.R., Hopkins, R.O., Reed, J.M., Kitchener, E.G., Squire, L.R., Carolina, N., Diego, S., 2003. Recognition memory and the human hippocampus. Neuron 37, 171–180.

Mayes, A.R., 1995. Memory and amnesia. Behav. Brain Res. 66 (1–2), 29–36.

Milner, B., 1966. Amnesia following operation on the temporal lobes. In: Whitiy, C. W.M., Zangwill, O.L. (Eds.), Amnesia. Butterworths, London, pp. 109–133.

Moran, R., Goshen-Gottstein, Y., 2015. Old processes, new perspectives: familiarity is correlated with (not independent of) recollection and is more (not equally) variable for targets than for lures. Cogn. Psychol. 79, 40–67.

Moscovitch, M., 1982. Multiple dissociations of function in amnesia. In: Cermak, I.L. S. (Ed.), Human Memory and Amnesia. Erlbaum, Hillsdale, NJ, pp. 337–370.

Moscovitch, M., 1994. Cognitive resources and dual-task interference effects at retrieval in normal people the role of the frontal lobes and medial temporal cortex. Neuropsychology 8 (4), 524–534.

Nadel, L., Moscovitch, M., 1997. Memory consolidation, retrograde amnesia and the hippocampal complex. Curr. Opin. Neurobiol. 7, 217–227.

Norman, K.A., O'Reilly, R.C., 2003. Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. Psychol. Rev. 110 (4), 611–646. http://dx.doi.org/10.1037/0033-295X.110.4.611.

Onyper, S.V., Zhang, Y.X., Howard, M.W., 2010. Some-or-none recollection: evidence from item and source memory. J. Exp. 139 (2), 341–364. http://dx.doi.org/10.1037/a0018926.Some-or-none.

Rajaram, S., 1993. Remembering and knowing: two means of access to the personal past. Mem. Cognit. 21 (1), 89–102 http://psycnet.apa.org/psycinfo/1993-40186-001.

Ratcliff, R., Sheu, C.F., Gronlund, S.D., 1992. Testing global memory models using ROC curves. Psychol. Rev. 99 (3), 518–535.

Rosenbaum, R.S., Gilboa, A., Levine, B., Winocur, G., Moscovitch, M., 2009. Amnesia as an impairment of detail generation and binding: evidence from personal, fictional, and semantic narratives in K.C. Neuropsychologia 47, 2181–2187.

Rosenbaum, R.S., McKinnon, M.C., Levine, B., Moscovitch, M., 2004. Visual imagery deficits, impaired strategic retrieval, or memory loss: disentangling the nature of an amnesic person's autobiographical memory deficit. Neuropsychologia 42 (12), 1619–1635.

Rosenstreich, E., Goshen-Gottstein, Y., 2015. Recollection-based retrieval is influenced by contextual variation at encoding but not at retrieval. PLoS One 10 (7), e0130403. http://dx.doi.org/10.1371/journal.pone.0130403.

Rotello, C.M., Heit, E., Dubé, C., 2015. When more data steer us wrong: replications with the wrong dependent measure perpetuate erroneous conclusions. Psychon. Bull. Rev. 22, 944–954.

Rotello, C.M., Macmillan, N.A., Reeder, J.A., 2004. Sum-difference theory of remembering and knowing: a two-dimensional signal-detection model. Psychol. Rev. 111 (3), 588–616. http://dx.doi.org/10.1037/0033-295X.111.3.588.

Rotello, C.M., Macmillan, N.A., Reeder, J.A., Wong, M., 2005. The remember response: subject to bias, graded, and not a process-pure indicator of recollection. Psychon. Bull. Rev. 12 (5), 865–873 http://www.ncbi.nlm.nih.gov/pubmed/16524003.

Rotello, C.M., Masson, M.E.J., Verde, M.F., 2008. Type I error rates and power analyses for single-point sensitivity measures. Percept. Psychophys. 70, 389–401.

Rotello, C.M., Zeng, M., 2008. Analysis of RT distributions in the remember-know paradigm. Psychon. Bull. Rev. 15 (4), 825–832. http://dx.doi.org/10.3758/PBR.15.4.825.

Rugg, M.D., Curran, T., 2007. Event-related potentials and recognition memory. Trends Cogn. Sci. 11 (6), 251–257. http://dx.doi.org/10.1016/j.tics.2007.04.004.

Rugg, M.D., Vilberg, K.L., 2013. Brain networks underlying episodic memory retrieval. Curr. Opin. Neurobiol. 23 (2), 255–260.

Rugg, M.D., Vilberg, K.L., Mattson, J.T., Yu, S.S., Johnson, J.D., Suzuki, M., Suzuki, M., 2012. Item memory, context memory and the hippocampus: fMRI evidence. Neuropsychologia 50 (13), 3070–3079. http://dx.doi.org/10.1016/j.neuropsychologia.2012.06.004.

Rugg, M.D., Yonelinas, A.P., 2003. Human recognition memory: a cognitive neuroscience perspective. Trends Cogn. Sci. 7 (7), 313–319. http://dx.doi.org/10.1016/S1364-6613(03)00131-1.

Ryan, J.D., Althoff, R.R., Whitlow, S., Cohen, N.J., 2000. Amnesia is a deficit in relational memory. Psychol. Sci. 11 (6), 454–461.

Schacter, D.L., Harbluk, J.L., McLachlan, D.R., 1984. Retrieval without recollection: an experimental analysis of source amnesia. J. Verbal Learn. Verbal Behav. 23, 593–611.

Schacter, D.L., Verfaellie, M., Anes, M.D., 1997. Illusory memories in amnesic patients: conceptual and perceptual false recognition. Neuropsychology 11, 331–342.

Schacter, D.L., Verfaellie, M., Pradere, D., 1996. Neuropsychology of memory illusions—false recall and recognition in amnesic patients. J. Mem. Lang. *35*, 319–334.

Slotnick, S.D., Dodson, C.S., 2005. Support for a continuous (single-process) model of recognition memory and source memory. Mem. Cognit. 33 (1), 151–170. http://dx.doi.org/10.3758/BF03195305.

Smith, C.N., Wixted, J.T., Squire, L.R., 2011. The hippocampus supports both recollection and familiarity when memories are strong. J. Neurosci. 31, 15693–15702.

Song, Z., Jeneson, A., Squire, L.R., 2011. Medial temporal lobe function and recognition memory: a novel approach to separating the contribution of recollection and familiarity. J. Neurosci. 31, 16026–16032.

Squire, L.R., 1992a. Declarative and nondeclarative memory: multiple brain systems supporting learning and memory. J. Cogn. Neurosci. 4, 232–243.

Squire, L.R., 1992b. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. Psychol. Rev. 99, 195–231.

Squire, L.R., Wixted, J.T., Clark, R.E., 2007. Recognition memory and the medial temporal lobe: a new perspective. Nat. Rev. Neurosci. 8 (11), 872–883. http://dx.doi.org/10.1038/nrn2154.

Swets, J.A., Green, D.M., 1963. Signal Detection by Human Observers. Massachusetts Institution of Tech. Cambridge Research Lab of Electronics.

Talmi, D., Goshen-Gottstein, Y., 2006. The long-term recency effect in recognition memory. Memory 14 (4), 424–436. http://dx.doi.org/10.1080/09658210500426623.

Tousignant, C., Bodner, G.E., 2012. Test context affects recollection and familiarity ratings: implications for measuring recognition experiences. Conscious. Cognit. 21, 994–1000.

Tulving, E., 1985. Memory and consciousness. Can. Psychol. 26 (1), 1–12. http://dx.doi.org/10.1037/h0080017.

Turriziani, P., Serra, L., Fadda, L., Caltagirone, C., Carlesimo, G.A., 2008. Recollection and familiarity in hippocampal amnesia. Hippocampus 18, 469–480.

Verfaellie, M., 1994. A re-examination of recognition memory in amnesia: reply to Roediger and McDermott. Neuropsychology 8 (2), 289–292.

Verfaellie, M., Treadwell, J.R., 1993. Status of recognition memory in amnesia. Neuropsychology 7 (1), 5–13.

Viskontas, I.V., McAndrews, M.P., Moscovitch, M., 2000. Remote episodic memory deficits in patients with unilateral temporal lobe epilepsy and excisions. J. Neurosci. 20 (15), 5853–5857.

Wais, P.E., Wixted, J.T., Hopkins, R.O.R.O.O., Squire, L.R., 2006. The hippocampus supports both the recollection and the familiarity components of recognition memory. Neuron 49 (3), 459–466. http://dx.doi.org/10.1016/j.neuron.2005.12.020.

Wickelgren, W.A., 1968. Sparing of short-term memory in an amnesic patient: implications for strength theory of memory. Neuropsychologia 6, 235–244.

Wixted, J.T., 2007. Dual-process theory and signal-detection theory of recognition memory. Psychol. Rev. 114 (1), 152–176. http://dx.doi.org/10.1037/0033-295X.114.1.152.

Wixted, J.T., 2009. Remember/Know judgments in cognitive neuroscience: an illustration of the underrepresented point of view. Learn. Mem. 16 (7), 406–412 http://learnmem.cshlp.org/cgi/content/abstract/16/7/406.

Wixted, J.T., Mickes, L., 2010. A continuous dual-process model of remember/know judgments. Psychol. Rev. 117 (4), 1025–1054. http://dx.doi.org/10.1037/a0020874.

Wixted, J.T., Stretch, V., 2004. In defense of the signal detection interpretation of remember/know judgments. Psychon. Bull. Rev. 11 (4), 616–641 http://www.ncbi.nlm.nih.gov/pubmed/15581116.

Woodruff, C.C. d, Hayama, H.R., Rugg, M.D., 2006. Electrophysiological dissociation of the neural correlates of recollection and familiarity. Brain Res. 1100 (1), 125–135. http://dx.doi.org/10.1016/j.brainres.2006.05.019.

Yonelinas, A.P., 1994. Receiver-operating characteristics in recognition memory: evidence for a dual-process model. J. Exp. Psychol.: Learn. Mem. Cognit. 20 (6), 1341–1354 http://www.ncbi.nlm.nih.gov/pubmed/7983467.

Yonelinas, A.P., 2001. Consciousness, control, and confidence: the 3 Cs of recognition memory. J. Exp. Psychol.: Gen. 130 (3), 361–379. http://dx.doi.org/10.1037//0096-3445.130.3.361.

Yonelinas, A.P., 2002. The nature of recollection and familiarity: a review of 30 years of research. J. Mem. Lang. 46 (3), 441–517. http://dx.doi.org/10.1006/jmla.2002.2864.

Yonelinas, A.P., Jacoby, L.L., 1995. The relation between remembering and knowing as bases for recognition: effects of size congruency. J. Mem. Lang. 34 (5), 622–643 http://www.sciencedirect.com/science/article/pii/S0749596 × 85710285.

Yonelinas, A.P., Kroll, N.E.A., Dobbins, I.G., Lazzara, M.M., Knight, R.T., 1998. Recollection and familiarity deficits in amnesia: convergence of remember-know, process dissociation, and receiver operating characteristic data. Neuropsychology 12, 323–339 http://bic.berkeley.edu/despolab/publications/pdfs/PDF.pdf.

Yonelinas, A.P., Kroll, N.E.A., Quamme, J.R., Lazzara, M.M., Sauvé, M.-J., Widaman, K. F., Knight, R.T., 2002. Effects of extensive temporal lobe damage or mild hypoxia on recollection and familiarity. Nat. Neurosci. 5 (11), 1236–1241. http://dx.doi.org/10.1038/nn961.

Yonelinas, A.P., Otten, L.J., Shaw, K.N., Rugg, M.D., 2005. Separating the brain regions involved in recollection and familiarity in recognition memory. J. Neurosci. 25 (11), 3002–3008. http://dx.doi.org/10.1523/JNEUROSCI.5295-04.2005.

Yonelinas, A.P., Parks, C.M., 2007. Receiver operating characteristics (ROCs) in recognition memory: a review. Psychol. Bull. 133 (5), 800–832. http://dx.doi.org/10.1037/0033-2909.133.5.800.

Yovel, G., Paller, K.A., 2004. The neural basis of the butcher-on-the-bus phenomenon: when a face seems familiar but is not remembered. Neuroimage 21 (2), 789–800. http://dx.doi.org/10.1016/j.neuroimage.2003.09.034.