Rapid Induction of False Memory for Pictures

Yana Weinstein

University College London

March 2009

Declaration

I, Yana Weinstein, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis. Experiment 3 and Experiments 5-7 appear in Weinstein and Shanks (2008) as Experiment 4 and Experiments 1-3 respectively.

Acknowledgments

My interest in pursuing an academic career was sparked when I took a casual research assistant job at Warwick University supporting Evan Heit in his research on reasoning. Although I never quite got the hang of that topic, working with Evan gave me a hint of the excitement involved in academic research. The following year, Kimberley Wade joined the department and I was overwhelmed by her lectures on memory and the law, and immediately identified my passion for false memory. When Evan suggested I apply to study with David Shanks, I was initially concerned that he would not find my ideas intriguing or relevant. I am glad to report that this did not turn out to be the case, and extremely indebted to David for his time and patience in advising me on a topic that was not necessarily one that he had always dreamt of delving into. I am also grateful to Leun Otten for her supporting role as secondary supervisor. I am very fortunate that University College London gave me this opportunity to learn how to conduct research, and that the Biotechnology and Biological Sciences Research Council made this opportunity a practical reality. Finally, I am also extremely grateful to the Bogue Fellowship which funded a trip to Washington University in St. Louis in 2006, where I had the honour of studying with Henry L. Roediger III whose name I remember from my very first class as a Psychology undergraduate. As luck would have it, Roddy invited me back. Thanks to family and friends for forgiving my fleeing to America.

Abstract

In this thesis, a new procedure is proposed which achieves the rapid induction of false recognition memory for rich pictorial stimuli. Chapter 2 presents the basic three-step procedure in which participants study some pictures, imagine others in response to words, and perform a picture recognition test. Imagining pictures leads to a false alarm rate of 27% above baseline (Experiments 1a and 1b). In Chapter 3, a source monitoring test is used to demonstrate that this effect is not solely due to diffuse familiarity (Experiment 2) and does not appear to be driven by perceptual processing (Experiment 3). Chapter 4 investigates whether imagination impairs discrimination between studied and new items presented side by side. On a two-alternative forced choice test, participants are less accurate in indicating the studied picture when it is paired with an imagined picture than when paired with a new picture (Experiment 4). The role of indirect tests in verbal false memory research and the value of applying them to pictorial stimuli is discussed in Chapter 5. An indirect perceptual identification test previously used for words is adapted to pictures and shown to be highly perceptually driven (Experiment 5). Chapter 6 presents consistent evidence for the lack of perceptual priming of imagined items in the current procedure despite high levels of false recognition (Experiments 6-7), again indicating that the effect is conceptually driven. Chapter 7 discusses how the new procedure presented in the thesis can further our understanding of the similarities and differences between true and false memories.

Table of Contents

List of Figures

List of Appendices

CHAPTER 1: FALSE MEMORY FOR VERBAL AND PICTORIAL STIMULI

1.1 False Memory: An Introduction

A feature of human memory is its susceptibility to distortions and illusions (Roediger, 1996). In the past 15 years, this susceptibility has generated much research into what is broadly defined as "false memory," although the use of this term has recently attracted some controversy (see DePrince, Allard, & Freyd, 2004; Pezdek & Lam, 2007; Wade et al., 2007). For our purposes, it suffices to say that this term refers to situations in which a memory does not correspond to the reality of an experience. The most straightforward example of a mismatch between memory and reality is when a person remembers experiencing an event which did not occur. Various methodologies have been used to study this phenomenon, ranging from elegant single-session experiments which elicit false recognition of prototypical dot patterns (Metcalfe and Fisher, 1986) and verbal stimuli (see Gallo, 2006 for a review), to long-term false memory inductions which lead participants to believe that they have experienced fictitious events (e.g., Loftus & Pickrell, 1995).

1.2 False Memory for Verbal Stimuli: The DRM Paradigm

The methodology most commonly used for studying memory errors for words is the Deese (1959) / Roediger and McDermott (1995) paradigm (known as the DRM paradigm) in which participants learn 15-word lists, each of which converges on one associate, the critical lure. For instance, the following list is based around the critical lure *CHAIR: table, sit, legs, seat, couch, desk, recliner, sofa, wood, cushion, swivel, stool, sitting, rocking, bench.* The lists used in this paradigm were simply designed by picking out the first 15 associates of the critical lure from a set of word norms (Russell & Jenkins, 1954). At test, despite never having seen it in the study phase, subjects spontaneously produce the critical lure *CHAIR* in free recall and/or report having studied it on recognition tests. While the lists vary in terms of how effectively they induce participants to acquire a false memory for the nonpresented critical lure, levels of false recognition (72%) can be higher than average true recognition (65%). In order to calculate corrected true and false recognition rates, baseline false recognition rates for items from non-studied lists (list items and critical lures respectively) were subtracted from the above figures. False recognition of critical lures from non-studied lists was at 16%, producing a corrected false recognition rate of 56%; analogously, false recognition of list words from non-studied lists was 11%, producing a corrected true recognition rate of 54%. More evidence for the strength of this false memory is that the critical lure was spontaneously produced in a free recall test at rate of 55%, which came close to the 62% mean probability of recall of the studied words (Roediger & McDermott, 1995).

1.3 False Memory for Autobiographical Events

While this paradigm has generated a lot of research (see Gallo, 2006), it has often been criticised for its lack of ecological validity, and some authors have pointed out the dangers of extending conclusions from research with such derived stimuli to controversial real-world situations such as recovered memories of childhood sexual abuse (Freyd & Gleaves, 1996). At the other end of the spectrum are studies which have looked at false memories for autobiographical events. In Loftus and Pickrell’s (1995) sample, 25% of participants came to falsely believe that they had got lost in a shopping mall at the age of five after a sequence of sessions spent trying to recall details from this event alongside details of events that had really taken place. Although this finding is striking and certainly highly applicable to everyday experience, the problem with this method is that *recognition* memory can never truly be tested. In experiments which use words and other derived stimuli, we can present the stimuli to participants and ask them whether a particular stimulus had been presented previously, whereas when studying autobiographical memory, we must rely on participants' introspective reports. Recognition is particularly important in the case of false memory because presentation of the stimulus itself at test allows participants to compare it directly to their memories of stimuli they had actually perceived, and to reject unstudied stimuli based on characteristics that appear novel.

1.4 False Memory for Pictures

Pictorial stimuli may be seen as offering a compromise between the ecological validity of autobiographical studies and the flexibility of verbal stimuli. However, there has been little research into false memory for pictorial stimuli, and when it has been conducted more often than not the stimuli have been line drawings (most commonly those developed by Snodgrass & Vanderwart, 1980). In studies looking into false memory for pictures, pictorial stimuli are either presented at encoding only, or both at encoding and at test. Of primary interest is the latter case, as such a procedure allows participants to compare the test stimulus directly to any existing memory trace and make a memory judgment on the stimulus itself.

In the *category associates procedure* (Hintzman, 1988; Koutstaal & Schacter, 1997), participants study lists of pictures which correspond to certain categories (e.g., animals) and are later given a recognition test consisting of some items shown at study, some novel items from the same semantic categories as the studied pictures (e.g., other animals), and some novel items from unstudied categories (e.g., clothes). This procedure typically produces impressive rates of corrected false recognition, because participants show minimal false recognition for items from unstudied categories and much higher levels of false recognition for new exemplars from studied categories (e.g., corrected false recognition of up to 40%, Koutstaal, Verfaellie, & Schacter, 2001). As the number of studied items within a category increases, so does the level of false alarms to new members of those categories. In these studies, high levels of false recognition arise because participants misapply the knowledge that they have studied many pictures from one category and incorrectly assume that any pictures from that category have probably been studied.

Israel and Schacter (1997) developed a version of the DRM paradigm in which each word was presented with an accompanying line drawing at study and/or at test. Studying and being tested on pictures lead to much lower levels of false recognition (corrected false recognition of 23%) than typically found on the standard word task. Finally, Fazendeiro, Winkielman, Luo, and Lorah (2005, Experiment 1B) used a similar but less powerful false memory induction procedure in which just one studied item was paired with a semantically related lure that was later presented at test along with studied items and unrelated lures, which produced a corrected false recognition rate of just 8% .

1.5 Picture Superiority and Distinctiveness

The finding that old and new pictures are easier to tell apart than old and new words is well-documented (Shepard, 1967). Various accounts have been proposed to explain this picture superiority effect. In his dual-coding theory, Paivio proposed that words and pictures are represented by individual codes and that the image code might be inherently stronger than the verbal code. Alternatively, Paivio suggested there could be cross-talk between these two distinct codes, but that the verbal code is triggered more easily by a picture rather than vice versa. So studying a picture might result in a memory of a picture plus an associated word generated internally, more often than studying a word would result in the memory of this word being accompanied by an internally generated image.

Another set of explanations does not call for two separate codes, instead attributing the picture superiority effect to the inherent distinctiveness of pictorial when compared to verbal representations. Pictures are more heterogeneous than words in terms of perceptual features, and thus have more distinctive features, which may facilitate the creation of unique memory traces (Nelson, Reed & McEvoy, 1977). In a similar vein, Nelson, Reed and Walling (1976) designed a set of line drawings that were so perceptually similar that they were no longer recognised with any more accuracy than words. The distinctiveness could also arise from conceptual processes, as pictures may offer more direct access to meaning than words, for instance by facilitating semantic tasks such as categorisation (Potter and Faulconer, 1975). The distinctiveness explanation has been applied to false memory research, although the focus has been on studies in which word/picture format is manipulated only at encoding and memory is tested with verbal cues. The idea behind the distinctiveness heuristic (Schacter, Israel, & Racine, 1999; Hege & Dodson, 2004) is that, having encoded items in a more distinctive format, participants search their memory for associated detail, and the absence of this detail can help them make a correct rejection.

The role of distinctiveness in picture recognition tests where study/test formats are matched is even more intuitive: in the case of studied items, the more distinctive features a picture has, the more likely it is that one of these features would have been encoded and can serve as a recognition cue (Loftus & Bell, 1975); while in the case of new items, the more distinctive details there are, the more likely it is that the novelty of one such detail will alert the participant to the fact that they had not seen the item before. Loftus and Bell compared recognition performance for black and white photographs and line drawings at exposure durations from 50 to 500 msec. Participants showed higher sensitivity to photographs compared to drawings at every exposure duration. Part of this effect was attributed to the increased likelihood of encoding an informative detail for a photograph than for a line drawing. However, the improved memory performance for photographs persisted even when participants claimed to be making the recognition response based on familiarity, implying that photographs are inherently more distinguishable than line drawings.

1.6 Monitoring and Misattribution

The most widely accepted explanation of the false memory phenomenon in the DRM paradigm involves a two-step process. First, the critical lure is activated by its semantic associates, i.e., the list of words that are related to this lure. This activation can occur both at encoding, when associates are actually studied, and at retrieval. Then, at test, a decision has to be made as to the source of this activation: the participant must decide whether they thought of this word themselves, or whether it had actually appeared in the list, a process described by Johnson, Hashtroudi, and Lindsay (1993) as *reality monitoring*. The activation/monitoring account explains false memory in the DRM paradigm as a combination of semantic activation and faulty reality monitoring which causes subjects to incorrectly attribute the activation to prior perception (Roediger, Watson, McDermott, & Gallo, 2001).

In fact, errors of misattribution are a very important concept in false memory. When people engage in reality monitoring, they rely on metacognitive cues to decide whether a memory is of an internally generated or an externally experienced event. Sometimes, these cues will be helpful and will generally lead to the correct decision. For instance, memories of externally experienced events tend to be richer in perceptual detail than internally generated content (Johnson, Suengas, Foley, & Raye, 1988); the latter, on the other hand, are commonly remembered together with related cognitive operations (Finke, Johnson, & Shyi, 1988; Schooler, Gerhard, & Loftus, 1986). Similarly, when trying to distinguish between true and false memories, participants may rely on the knowledge that they tend to recall more distinctive, item-specific perceptual details for true than false memories (Norman & Schacter, 1997). Other such cues can be misleading. A classic example is the *fluency heuristic*: if some material is processed more fluently than other material on a test, or simply more fluently than expected, this material evokes a feeling of familiarity and this familiarity is then interpreted as arising from prior study (Johnston, Dark, & Jacoby, 1985). However, this may not always be a correct inference: fluency can be affected by a variety of factors other than prior study. In the verbal domain, for instance, perceptual clarity

(Whittlesea, Jacoby, & Girard, 1990) and meaningful semantic context (Whittlesea, 1993) can make words easier to process and hence appear more familiar.

Fazendeiro et al. (2005; Experiment 3) demonstrated that misattributed fluency can contribute to false recognition even when formats are not matched between study and test. Participants studied a mixed list of line drawings and words, and were tested on studied items, semantically related items, and new items. Studied items appeared in the same format (word or line drawing) as at study, while related items were presented in the opposite format. So, for instance, the word *chair* may have been presented in the study phase, and a picture of a table at test. Corrected false recognition was .13 for pictures whose semantic associates had been presented as words in the study phase. As with the DRM paradigm, this result could be explained in terms of activation followed by faulty monitoring, with the caveat that any such activation must be exclusively conceptual. The current procedure makes use of the transfer of conceptual fluency across presentation format and induces false recognition of pictures by presenting participants with the names of the objects depicted in those pictures.

1.7 False Memory for Pictures: A New Procedure

Above we reviewed evidence that people do better in memory tests when the stimuli are more visually complex, be it pictures as opposed to words, or highly differentiated line drawings as opposed to ones that are similarly structured. It is perhaps because of this picture superiority phenomenon that researchers have tended to shy away from false memory procedures where pictures are presented both at study and at test. The first goal of this thesis was thus to develop a false memory induction using distinctive photographs, without relying on long delays between study and test or driving down correct recognition of studied items. In order to achieve this, the current experiments involve a three-step procedure based on misinformation (Loftus, 1975) and imagination inflation (Johnson, Raye, Foley, & Kim, 1982) studies. Chapter 2 provides a description of how these and other false memory experiments inspired the design of the current procedure, and describes two versions (Experiments 1a and 1b) which demonstrate its success in inducing high levels of false recognition of distinctive photographs.

The first two experiments employed variations on the standard old/new recognition test. It was important to make sure that false recognition was not attributable to factors other than a genuine memory of having perceived the picture. In order to verify this, the testing procedure was modified to make it more sensitive, first using a source monitoring test (Chapter 3, Experiments 2 and 3) and then a two-alternative forced choice test (Chapter 4, Experiment 4). Although these tests were more conservative, they produced significant levels of false recognition which were hard to explain without invoking the existence of a genuine false memory for the unstudied pictures.

1.8 False Memory on Indirect Tests

Another way of testing memory is through indirect or implicit tests, whereby prior exposure to test stimuli facilitates performance on tasks including lexical decision (Ratcliff, Hockley, & McKoon, 1985) and picture naming (Warren & Morton, 1982). Recently, researchers have investigated whether false memory for critical lures in the DRM paradigm is exhibited on indirect as well as direct memory tests. Indirect tests can be roughly divided into those that tap into the perceptual and those that tap into the conceptual details of studied stimuli (Weldon, Roediger, Beitel, & Johnston, 1995). Perceptual indirect tests such as speeded recognition of degraded stimuli rely on the identification of perceptual features, while conceptual indirect tests such as word association access an item's meaning. It has generally been found that performance on conceptual indirect tests is facilitated for critical lures as well as for studied items (e.g., McDermott, 1997), while the behaviour of critical lures on perceptual indirect tests is more controversial (e.g., Hicks & Starns, 2005; McBride, Coane, & Raulerson, 2006); these findings are reviewed in Chapter 5. So far, no work has been done on false memory for pictorial stimuli using indirect memory tests. In Chapter 6 (Experiments 5–7), the indirect test methodology is applied to pictures in order to investigate whether false memory for pictures in the new procedure presented in Chapter 2 is driven by perceptual processes.

1.9 Perceptual Processes in False Recognition of Pictures

As well as developing a new procedure that produces high levels of false recognition of pictures at short delay, another aim of this thesis is to specify the nature of the representations underlying the false memories created with the new experimental procedure. An important question is whether false memories can possess similar perceptual attributes to true memories. Attempts have been made to answer this question by examining participants' introspections regarding the perceptual qualities of true and false memories, either with verbal reports (e.g., Lampinen, Odergard, & Bullington, 2003; Heaps & Nash, 2001; Norman and Schacter, 1997; Schooler et al., 1986) or rating scales (e.g., Mather, Henkel, & Johnson, 1997; Johnson et al., 1988). Although the results have been mixed, the phenomenological differences between true and false memories appear to be quite subtle.

One suggested explanation for the blurred distinction between true and false memories is that false memories may acquire perceptual details by borrowing them from true memories (e.g., Lampinen, Meier, Arnal, & Leding, 2005). However, false memories do not need to be perceptually based. A false memory might arise when an internally generated perceptual representation is mistaken for a memory of perceiving the item, but it could also arise from nonperceptual activation. This question is addressed throughout the thesis (in particular in Experiments 3 and 5-7), and various techniques are used to determine whether the false memories reported by participants in the current procedure are conceptually or perceptually based.

CHAPTER 2: FALSE MEMORY FOR PICTURES: A NEW PROCEDURE

One process that has been found to contribute significantly to the formation of false memories is that of imagination. While the benefits of imagination to memory enhancement are well known (e.g., Marschark & Surian, 1989), imagination has also been used as a tool to induce false memories of actions (Anderson, 1984; Goff & Roediger, 1998), auditory and visual events (Henkel, Franklin, & Johnson, 2000), autobiographical memories (Garry, Manning, Loftus, & Sherman, 1996; Hyman, Husband, & Billing, 1995; Hyman & Pentland, 1996), words (Johnson, Raye, Wang, & Taylor, 1979), pictures of visual scenes (Intraub & Hoffman, 1992), and pictures of objects (Henkel & Franklin, 1998; Johnson et al, 1979).

According to the source monitoring framework, imagined and perceived events differ in terms of their memory characteristics, and these characteristics are useful when people are trying to tell the two apart (Johnson & Raye, 1981). For instance, memories of perceived events are thought to be rich in sensory detail (e.g., Stern & Rotello, 2000), so when deciding whether an event had been perceived or imagined, people might set a criterion for the level of sensory detail recalled. However, there are strong suggestions from neuroimaging work that imagination and perception overlap both in terms of function (Kosslyn & Thompson, 2000) and structure (Borst & Kosslyn, 2008), and there is some evidence that false alarms are predicted by activation of brain regions linked to vivid imagination (Gonsalves and Paller, 2000; Gonsalves, Reber, Gitelman, Parrish, et al., 2004). More specifically, Gonsalves and colleagues postulate that the likelihood of an imagined event being misattributed to perception depends on the vividness of the imagery.

If imagination can produce activation that is similar in nature to that resulting from perception, that would make it very difficult to differentiate between perceived and imagined events. In the current procedure, imagination inflation is used in an attempt to create an internal event that can be mistaken for a true memory.

A method of inducing false memory made popular by Loftus (1975) involves the insertion of misleading information between study and test. Loftus, Miller, and Burns (1978) showed participants a set of slides depicting an event and later asked participants to choose the correct slide from pairs in which one slide was identical to the original and the other was slightly altered. In between the study and test phases, participants were either given verbal information consistent with the original slides, or consistent with the lure slides. In the latter condition, participants were more likely to choose the lure slide than the slide they had actually seen. This effect appears to be striking, but is undermined by the fact that the two pictures in each pair were identical apart from the one detail referred to in the misleading information. Okado and Stark (2005) used a similar three-step procedure to inflate false memory for pictures. In the study phase, participants were presented with some object names accompanied by a picture and imagined other objects, while at test they were presented with the names of studied, imagined and new pictures and asked to indicate which ones had been presented with a picture. In order to inflate the number of false alarms made on imagined trials, in an intervening phase participants were presented with some of the object names from the study phase and encouraged to lie about having seen pictures they had only imagined. In a pilot test, this manipulation lead to an increase from 13% to 17% of false endorsements of imagined items as having been studied.

In the current procedure, a three-phase design is combined with the categorised list structure of the category associates procedure described in Chapter 1 in an attempt to elicit high levels of false recognition of imagined pictures. In the study phase, participants label pictures (black and white or colour photographs) of objects that belong to one of five categories and at test are presented with pictures and asked to indicate which ones had appeared in the study phase. The false memory induction occurs in an intervening phase where participants are presented with the names of the studied objects, with a small proportion of items replaced with names of objects that had not appeared in pictorial form in the study phase but belong to the same five semantic categories. Participants are not informed of this manipulation; the instructions for this phase require them to recall whether the object had appeared in black and white or in colour in the study phase. In the absence of such a memory (i.e., for those objects names that refer to non-presented pictures, although this is not made explicit), participants are asked to imagine the picture and answer the question based on this image. In the test phase, participants are presented with studied pictures, imagined pictures, and new pictures, all of which belonged to the same five semantic categories. Experiments 1a and 1b demonstrate that this imagery manipulation – asking people to imagine an object in order to make a recollective judgment about a picture they had never seen – is very successful in eliciting high levels of false recognition.

2.1 Experiment 1a

The test phase of this experiment consisted of an old/new recognition test with remember/know judgments. The distinction between remembering and knowing was first used experimentally by Tulving (1985) to tap into the subjective difference between memories that feel like facts (noetic memory, associated with *know* responses) and memories that can be re-experienced by the process of mental time travel (autonoetic memory, associated with *remember* responses). False alarms are more frequently made with *know* responses, with some exceptions (e.g., in the DRM paradigm, Roediger and McDermott, 1995). In the current procedure, the distribution of remember/know responses was compared for two types of false alarms: false alarms to new items and false alarms to previously imagined items.

Method

Participants

Twenty-four participants (8 males) aged 18 to 30 (M = 20.0) took part in the experiment. Participants were volunteers recruited within University College London and reimbursed for their time.

Materials

The stimuli in the experiment consisted of 110 categorised pictures normed for name agreement, and 24 additional pictures used as filler items. For the corpus of 110 categorised pictures, lists of 22 items were constructed for the following five categories: animals, clothes, electrical appliances, fruit and vegetables, and household objects (see Appendix A for full list of items). The selection of items in each list was based as far as possible on semantic category norms (Van Overschelde, Rawson, & Dunlosky, 2004), but constrained to some extent by availability of appropriate pictorial stimuli. In order to achieve the required list length, Van Overschelde et al.'s "fruit" and "vegetable" categories were merged to form one list, while "kitchen utensils" and "carpenter's tools" were combined to create the household objects list, and additionally a list of electrical appliances was compiled. There was no overlap between categories, so, for instance, none of the household objects could be powered by electricity. An additional 24 pictures were used as fillers in the practice task prior to the study phase (14 furniture items), and as fillers in the test phase (5 musical instruments and 5 food products).

Based on the semantic category norms, six critical items were chosen from each of the four categories for which norms were available. Each item was one of the top seven most popular items in each category. For the new category of electrical appliances, norms were produced by asking 20 people to list six exemplars; the six most popular exemplars were used as critical items for the electrical appliances category.

Figure 2.1: Sample stimuli used in all experiments. The left panel shows a full colour picture while the right panel shows another picture converted to black and white format.

The stimuli were full-colour photographs edited in Adobe Photoshop. Items were presented on a 300 x 300 pixel white background with shadows and extraneous detail such as text removed (see Figure 2.1 for sample stimuli). Black and white versions of the pictures were produced using the Grayscale function. In order to obtain agreed labels for the pictures, four people were asked to name a larger pool of pictures, half of which appeared in colour and half in black and white; this was counterbalanced between participants. The final set of 110 categorised pictures included only those which elicited the same response from at least three of the four participants. The study, imagery and test phases of the experiment were programmed in Visual Basic 6.0 and a PC with screen resolution of 1024 x 768 pixels was used to run the experiment.

Design

The experiment used a repeated measures design, with item type (seen/imagined/new) as the within-subjects variable. Of the 110 categorised pictures, a total of 30 were critical items (six from each of the five categories; see italicised items in Appendix A). There were three groups of critical items: 10 seen items, 10 imagined items, and 10 new items. Seen items were those that were presented to participants as pictures during the study phase and as labels in the imagery phase; imagined items were those that appeared only as incorrect options in the study phase labelling task and once again as labels in the imagery phase; and new items did not appear in either phase. The three sets of 10 critical items (two from each of the five categories in each set) were rotated between the seen, imagined and new conditions giving three counterbalancing conditions. Apart from the critical items, participants in all conditions also studied the remaining 80 non-critical items, i.e., 16 pictures each from the animals, clothes, electrical appliances, fruit and vegetables, and household objects categories. In each of these conditions, half the items appeared in colour and half in black and white in the study and test phases. A further three conditions were created to counterbalance this, producing six counterbalancing conditions in total.

Procedure

The experiment was completed in one 30-minute session and involved a study phase in which participants labelled a set of pictures, an imagery phase where participants tried to recall whether the pictures they had seen in the study phase had been presented in black and white or in colour, and a recognition test phase.

Study Phase. In this phase, participants were shown a total of 90 pictures, made up of 80 non-critical items common to all conditions and 10 pictures from the critical items list (these are referred to as *seen* pictures at test). Their task was to choose the correct word for each picture from a choice of two. The two words appeared on the screen first for 4000 msec, one either side of the screen. During this time participants were asked to imagine pictures for each of these words. They were not given any specific instructions for this imagery task (e.g., whether to imagine the pictures in black and white or colour). Next, a picture corresponding to one of the words appeared in the centre of the screen for 250 msec and the words disappeared. Following this brief presentation, participants had to make a choice by pressing either the \ key to select the left word or the / key to select the right word. As a reminder, these symbols appeared in place of the words when it was time to make a choice. As soon as participants made their selection, a low-pitched beep sounded in the case of an incorrect response, and the next trial began immediately.

The word corresponding to the picture appeared as one of the choices on each trial. The alternative word corresponded to a different picture from the same category. For example, participants may have been shown the words *cat* and *dog* followed by a picture of a dog. Out of the 18 trials of each category, on 12 trials the incorrect word corresponded to another picture participants would see during the experiment. For the other six trials in each category, incorrect words corresponded to the two critical imagined items which participants would not see in the study phase. The order of trials was randomised across all categories, and word pairs were chosen at random by the program with the above constraints.

Participants read instructions on screen and completed a practice block of 14 items, after which they were told their score and the real experiment began. Of the 90 pictures shown to participants, 45 appeared in colour and 45 in black and white, but there was nothing in the instructions relating to this. There was also no mention of a memory test and participants were not asked explicitly to try to memorise the pictures or anything about them.

Imagery Phase. Following a 5-minute mental arithmetic distractor task, participants were shown 90 words and asked to indicate whether they had seen the picture corresponding to this word in black and white or in colour during the first part of the experiment. This task was used in order to encourage participants to create and examine a mental image of the picture. Each word appeared for 4000 msec in the centre of the screen, during which time participants were asked to form a mental image of the corresponding picture in order to answer the question. Once the word disappeared, choices were made using the same keys as in the study phase. Participants were asked to use the number keypad to rate how confident they were in their answer following each response on a scale of 1 to 5 where 1 = *not at all confident* and 5 = *extremely confident*.

The words that appeared in this phase referred to 70 of the 80 pictures seen by all participants in the study phase (the same 10 pictures were excluded for all participants), and the items that later appeared as seen items (10 words) and imagined items (10 words) in the test phase. Eighty of the words thus corresponded to 80 of the pictures seen by participants in the study phase, while the other 10 had not previously been encountered as pictures. The instructions stated that the words in this phase corresponded to pictures seen in the previous part of the experiment. In other words, participants were not informed that some of the words had not actually appeared as pictures. However, they were instructed to try to form mental images for any words they could not remember having seen as pictures in the study phase and make their choices based on these mental images in the absence of a memory. The format of this phase was similar to that of the study phase, with instructions appearing on screen. There was no practice block and no feedback in response to participants' answers.

Test Phase. All participants, regardless of counterbalancing condition, were tested on the same 40 items. Ten of these were filler items from novel categories, and the other 30 were critical items: 10 seen items, 10 imagined items and 10 new items. Of the 40 pictures, participants had therefore only actually seen 10 during the study phase. Pictures were presented in the same central location as during the study phase. Seen pictures appeared in the same colour format as in the study phase, while new pictures appeared in black and white or colour according to counterbalancing condition, so that half of the items seen at test were in black and white and half were in colour. Participants were shown the 40 pictures one by one and asked to judge whether they were *old* (i.e., whether participants had seen them in the study phase) or *new* using the same two keys they had used in the study and imagery phases. Instructions appeared on screen as in the other phases. Participants made old/new judgments in their own time while viewing the picture. If participants identified a picture as old, the picture remained on screen and they were asked to make a remember/know judgment. Participants were asked to press R if they had a conscious recollection of the item's occurrence during the first stage, or of what they were thinking about at the time it was presented, and to press K if they were sure that the item was presented in the study phase, but were unable to consciously recollect its actual appearance or any other details regarding its occurrence in the study phase. Following a new or remember/know judgment, the next trial began immediately. There was no practice phase for this test.

Results

Study Phase

Performance in the study phase ranged from 93.3% to 100% correct identification of pictures from a choice of two labels, with a mean score of 97.8%. Over half the participants made less than three identification errors in the 90 trials, and no-one made more than eight errors.

Imagery Phase

Performance in the imagery phase was based on the 80 items which had appeared in the study phase. All participants performed above chance in this task, in which they had to indicate whether they had seen each corresponding picture in black and white or colour in the study phase. Scores ranged from 53.8% to 81.3%, with a mean score of 64.6% (SD = 7.0), which was significantly different from the chance level of 50%, $t(23) = 10.11$, $p < .001$. On the 1 to 5 confidence scale, participants were significantly more confident on their answers to seen items (*M* = 3.43) than to imagined items (*M* = 2.10); $t(23) = 8.40$, *SEM* = .16, $r = .74$, $p < .001$.

Test Phase

The left panel of Figure 2.2 shows the proportion of *old* responses to the three item types (seen, imagined, and new) by colour format (black and white or colour). Participants responded almost identically to colour and black and white pictures ($F < 1$), so this factor was left out of all further analyses. The responses shown in Figure 2.2 represent hits for seen items and false alarms for imagined and new items.

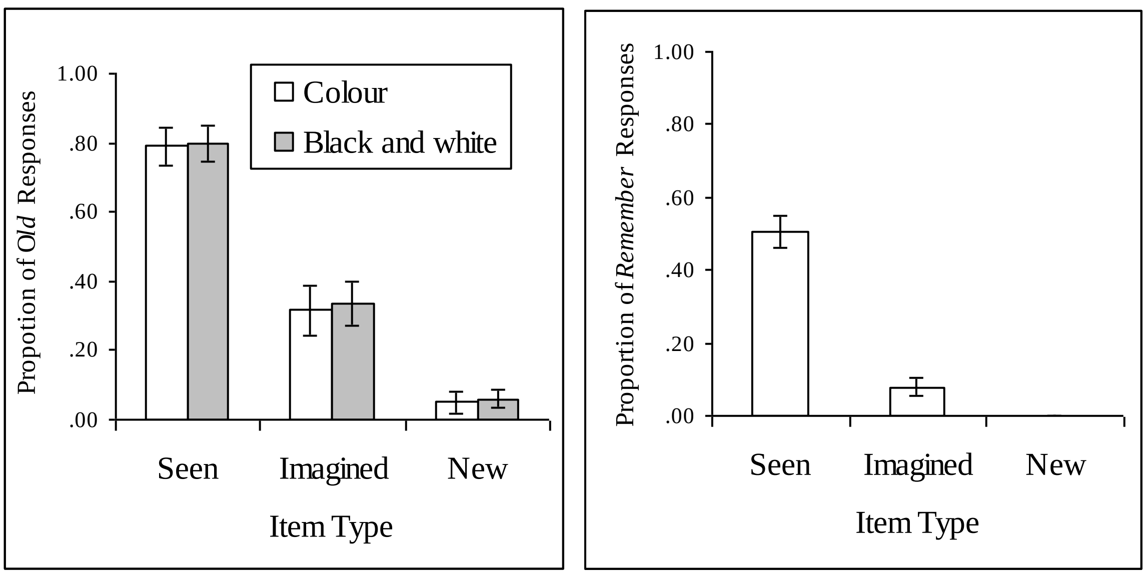

Figure 2.2: Left panel: Proportion of old *responses to seen, imagined, and new items by colour format in Experiment 1a. Right panel: Proportion of* remember *responses to seen, imagined, and new items across colour formats in Experiment 1a. Error bars represent SEM in this and all subsequent figures.*

A repeated measures ANOVA showed that there was a significant effect of item type, $F(1.61, 37.0) = 209.72$, $MSE = .02$, *partial* $\eta^2 = .90$, $p < .001$ (Mauchly's test indicated that the assumption of sphericity had been violated: $\chi^2(2) = 6.14$, $p < .05$; a Greenhouse-Geisser estimate of sphericity was used to correct degrees of freedom, $\varepsilon = .80$). The likelihood of participants responding *old* to an item was therefore dependent upon its status. The proportion of hits to pictures which had appeared in the study phase ($M = .79$) as well as false alarms to pictures which had appeared as words in the imagery phase ($M = .27$) were both significantly higher than the proportion of false alarms to new pictures ($M = .05$; $t(23) = 26.86$, $SEM = .03$, $r = .96$, $p < .001$ and $t(23) = 5.72$, $SEM = .04$, $r = .59$, $p < .001$ respectively), and 21/24 participants showed the latter effect (i.e., the number of false alarms they made to imagined items was at least one item higher than the number of false alarms they made to new items). The right panel of Figure 2.2 shows the proportion of *remember* responses by item type. *Remember* responses were made on 64% of hits to seen items (.50 of all seen trials) and on 29% of false alarms to imagined items (.08 of all imagined trials). Half of the participants made at least one *remember* false alarm on an imagined trial, whereas every false alarm to every new item across all participants was accompanied by a *know* response.

In this and all subsequent experiments, baseline estimates of bias are presented for seen versus new items. The bias results are summarised in Appendix B. *C* was calculated from hits and false alarm rates; to accommodate hit rates of 1 and false alarm rates of 0, 0.5 was added to each datapoint, which was then divided by N + 1, where N is the number of trials of that type (Snodgrass & Corwin, 1988). Given that one item type would have had to be kept constant between pairs, there was no additional value in comparing bias between any two pairs of items. Instead, the baseline bias estimate was compared to zero in all cases, and between conditions where there was more than one. In the current experiment, participants responded with a mean criterion of .37 which was significantly stricter than zero, $t(23) = 4.76$, $p < .001$.

Discussion

Experiment 1a successfully implemented a new procedure which elicited high levels of false memory for rich pictorial stimuli with a false memory induction involving imagination: imagining pictures led to an increase of .22 in false alarms compared to new items. These false alarms were sometimes accompanied by feelings of recollection relating to the encoding event: almost a third of the imagined item false alarms as compared with none of the new item false alarms were made with a *remember* response. The high hit rate and low new item false alarm rate indicates that participants had no problem distinguishing between items they previously studied and new items from the same category. The elevated false recognition rate for the imagined items was thus specifically related to studying the object's name and/or imagining the pictorial representation of this object.

Performance did not differ for colour compared to black and white pictures. Although the black and white/colour manipulation was only included in order to provide a memory task for the imagery phase, it also served as a useful test of the effects of colour information on memory. There was no difference in the way participants responded to colour and black and white stimuli in the present experiment. While some studies have shown that colour can be used to enhance memory for pictures (Bousfield, Esterson, & Whitmarsh, 1957; Suzuki and Takahashi, 1997), others have found no improvement in memory from the inclusion of colour information (Paivio, Rogers, and Smythe, 1968), and the results reported here support the latter finding. Note that the experiments presented here were not designed to explicitly test any hypotheses relating to stimulus distinctiveness and its effects on false memory; rather, detailed and varied pictorial stimuli were used in all experiments as a response to the findings already documented in the literature and reviewed in Chapter 1 regarding these effects.

2.2 Experiment 1b

Experiment 1a allowed participants only three response options (*remember*, *know*, and *new*). In the current experiment, a more sensitive recognition test was employed. Instead of making old/new and remember/know judgments, participants responded to each item on a 6-point confidence scale. This allowed the construction of ROC curves and established that false alarm rates to imagined and new items differed across all confidence levels.

Method

Participants

Thirty-six participants (20 males) aged 18 to 40 (*M* = 25.7) took part in the experiment. Participants were reimbursed for their time.

Procedure

The procedure of this experiment was identical to that of Experiment 1a except for response options in the test phase. Recognition judgments were made on a confidence scale from 1 to 6 where 1 = *very sure new* and 6 = *very sure old* to allow the construction of empirical ROC curves.

Results

Study and Imagery Phases

In the study phase, participants identified pictures from a choice of two words with a mean correct score of 97.2%. Performance in the imagery phase, where participants reported the colour format of pictures seen in the study phase, ranged from 55.3% to 81.3%, with a mean score of 65.0% (*SD* = 6.9), which was significantly above the chance level of 50%, $t(35) = 13.11$, $p < .001$. As in Experiment 1a, in the imagery phase participants were significantly more confident on the 5-point scale in their answers to the 80 pictures from the study phase (*M* = 3.54) than to the 10 novel items (*M* = 2.48), $t(35) = 9.66$, *SEM* = .11, *r* = .52, $p < .001$. There was no difference in performance on these measures between Experiments 1a and 1b.

Test Phase

Initially, responses given on the 6-point scale were collapsed into *old* responses (items given a rating of 4–6) and *new* responses (ratings of 1–3). The left panel of Figure 2.3 shows the proportion of *old* responses to each item type. A repeated measures ANOVA showed that there was a significant effect of item type, $F(2, 70) = 170.40$, *MSE* = .03, *partial* $\eta^2 = .83$, $p < .001$. The likelihood of participants responding *old* to an item was therefore dependent upon its status. Planned paired-samples t-tests showed a significant difference in *old* responses to seen items (*M* = .78) when compared to new items (*M* = .06; $t(35) = 20.64$, *SEM* = .04, *r* = .93, = $p < .001$) and also imagined items (*M* = .37) compared to new items, $t(35) = 7.05$, *SEM* = .04, *r* = .62, $p < .001$; participants made significantly more false alarms to imagined items than to new items, and 30 of the 36 participants showed this effect (i.e., they made at least one more false alarm to imagined items compared with new items).

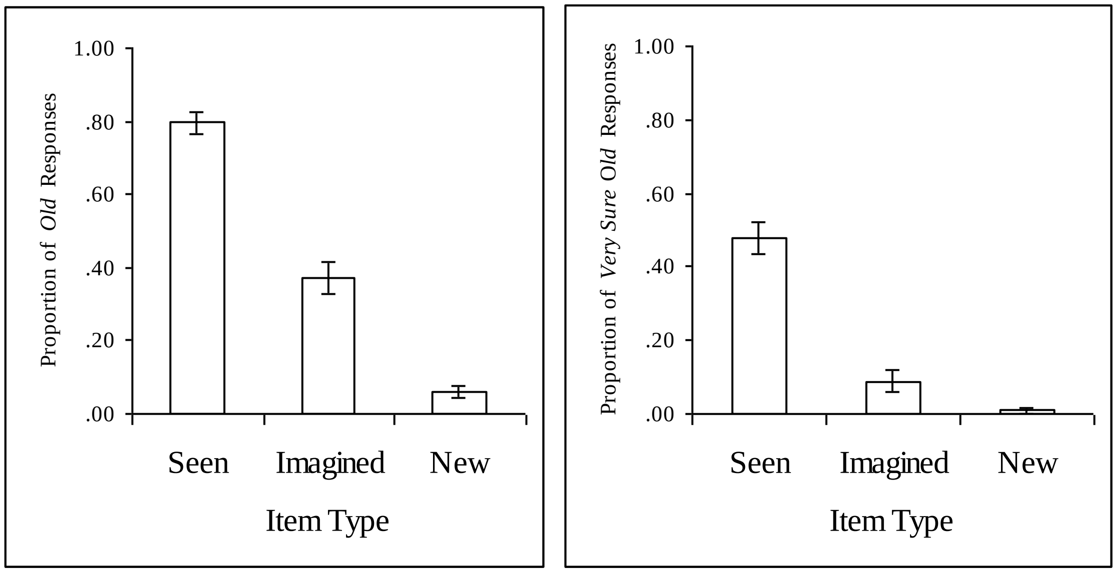

Figure 2.3: Left panel: Proportion of old *responses to seen, imagined, and new items in Experiment 1b. Right panel: Proportion of* very sure old *responses to seen, imagined, and new items in Experiment 1b.*

The right panel of Figure 2.3 shows the proportion of high confidence (confidence level 6, *very sure old*) responses by item type. The same pattern of data emerged when these responses were analysed separately: a repeated-measures ANOVA showed a significant effect of item type, $F(1.67, 58.6) = 89.41$, $MSE = .03$, *partial* $\eta^2 = .72$, $p < .001$ (Mauchly's test indicated that the assumption of sphericity had been violated: $\chi^2(2) = 7.3$, $p < .05$; a Greenhouse-Geisser estimate of sphericity was used to correct degrees of freedom, $\varepsilon = .84$). Even at a high confidence level, false alarms to imagined items were made at a mean rate of .09, which was significantly higher than the false alarm rate of .01 to new items, $t(35) = 2.84$; $SEM = .03$, $r = .30$, $p < .01$; only two participants made any high confidence false alarms to new items, while 15 of the 36 participants made high confidence false alarms to imagined items. High confidence hits were made at a rate of .48; the high confidence responses thus replicated the pattern obtained for *remember* responses in Experiment 1a. In fact, a post-hoc cross-experimental comparison indicated that the criterion set for high confidence responses in the seen-new comparison was .98, almost exactly the same as the criterion for responding *remember* in Experiment 1a (.97; $p > .9$).

Due to the low number of false alarms to new items, average confidence of *old* responses was compared between the seen and imagined conditions only for those 17 participants who made at least 4 *old* responses in each of the two categories. A within-subjects ANOVA confirmed that hits to seen items were made with higher confidence than false alarms to imagined tems ($M = 5.46$ and $M = 4.80$ respectively; $F(1, 16) = 46.21$, $MSE = .08$, *partial* $\eta^2 = .74$, $p < .001$). The confidence ratings collected in this experiment permitted ROC curves to be plotted for the seen-imagined and seen-new comparisons. ROC curves are shown in Figure 2.4 and were constructed by plotting the cumulative hit and false alarm rates at each confidence level, starting at confidence level 6. It is clear from Figure 2.4 that the points from the two comparisons fall on different curves. As in Experiment 1a, the criterion in the seen-new comparison, ($M = .30$; see Appendix B), was significantly stricter than zero, $t(35) = 4.44$, $p < .001$[1].

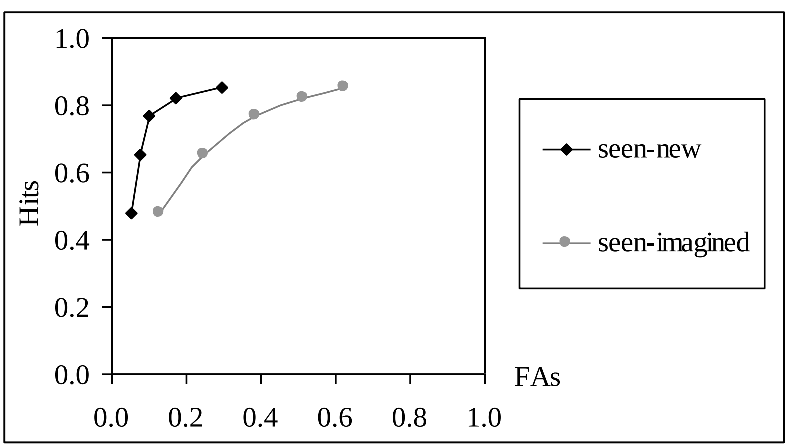

Figure 2.4: ROC curves for seen-new and seen-imagined comparisons in Experiment 1b.

Discussion

This experiment replicated and extended the effect presented in Experiment 1a. The current version of the procedure produced even higher corrected false recognition rates than the first. One explanation for this is that the effort of making an additional (remember/know) judgment following an *old* response may have encouraged participants to be more conservative on trials where they were not confident that they had seen the item, whereas in the current experiment equal effort was associated with responding *old* and *new*. Tentative evidence for this explanation is the numerical decrease in bias from Experiment 1a to 1b, although in a between-experiment comparison this difference did not reach significance.

In the current experiment, participants made more false alarms to imagined than new items at every confidence level. In fact, there were more false alarms to imagined items at the highest confidence level (*very sure old*, $M = .09$) than there were false alarms to new items at any confidence level (*very sure old*, *sure old*, and *guess old*, $M = .06$). Participants are thus clearly differentiating between imagined and new items on the test despite being asked whether they had seen the picture itself, and in some cases are very confident that they had previously studied a picture they had only imagined.

2.3 Discussion of Experiments 1a and 1b

Experiments 1a and 1b demonstrated a new procedure which produced corrected false recognition rates of up to 30% (Experiment 1b) for rich pictorial stimuli presented only 15 minutes earlier. Unlike most previous studies looking into false recognition at short delays (e.g., Israel & Schacter, 1997), the current procedure uses rich pictorial stimuli and always presents the pictures themselves as retrieval cues. The distinctiveness of the pictures should have aided participants in rejecting unseen stimuli, and yet the reported false recognition was higher than that in Israel and Schacter's pictorial version of the DRM paradigm. False recognition of imagined items was not quite as high as false recognition of pictures from studied categories reported in some versions of the category associates procedure (e.g., Koutstaal, Verfaellie, & Schacter, 2001). This is unsurprising given that in the category associates procedure, if participants relied only on their memory for which categories they had studied, they would automatically endorse any pictures from studied categories as old. In the current procedure, on the other hand, such a strategy would have increased not only false alarms to imagined items, but also false alarms to new items as both types of items were members of studied categories. On the contrary, the level of false recognition of new items in Experiments 1a and 1b was no higher than the false recognition rate typically obtained for items from unrelated categories in the category associates procedure. According to the fuzzy trace theory (Reyna & Brainerd, 1995), false alarms to new items from previously studied categories are made when participants are very reliant on gist information, whereas correct recognition would require a focus on item-specific information. In the present procedure, participants are clearly paying considerable attention to such information, which allows them to identify studied pictures and reject new pictures from studied categories with great accuracy.

One criticism of these experiments is that, in the recognition test, participants could simply be responding *old* to pictures when they are aware that they had only studied the associated word, either because they are misunderstanding the instructions or because they are relying on familiarity and endorsing any familiar item as *old*. A source-monitoring test (Johnson et al., 1993) would give participants the option to indicate that they had previously only imagined a picture (or encountered it in word form) and not studied the picture itself. Given this additional response option, if false alarms to imagined items were only made in Experiments 1a and 1b due to familiarity processes or misunderstanding of instruction, we would expect no more imagined than new pictures to be incorrectly classified as having been studied. Chapter 3 explores this possibility.

CHAPTER 3: FALSE MEMORY FOR PICTURES ON A SOURCE MONITORING TEST

Henkel and Franklin (1998) investigated imagination inflation with a source monitoring test. In their experiments, participants were shown drawings of various objects, and asked to imagine other objects. At test, participants were prompted with object labels and asked whether they had perceived or imagined the associated object; in other words, they were asked to identify the source of their memory for this item. One way in which participants might perform this task is by imagining (or re-imagining) a picture of the object represented by the label, and then comparing this picture to any existing memory of the item. It is not too hard to see how in this situation participants might have trouble identifying whether the memory is of a self-generated image or a previously presented picture. Indeed, Henkel and Franklin found that increasing the number of perceived objects that were similar to an imagined object made participants more likely to incorrectly identify the imagined object as having been perceived.

In Johnson et al.'s (1993) source monitoring framework, the first stage of source attribution is to decide whether a memory is of an internally generated event or an event perceptually experienced in the outside world. Various aspects of the memory might help to distinguish between these two categories – for instance, one might examine the vividness of the image. Incorrect source attribution in the case of an imagined picture could thus result from the incorrect use of a heuristic which attributes a vivid mental image to prior perception. Source monitoring failure can be seen to play out in real life situations, for instance in the mugshot exposure effect where eyewitnesses mistakenly pick an innocent member of a line-up because they recognise them from a mugshot rather than the original event (see Deffenbacher, Bornstein, & Penrod (2006) for review), and in autobiographical memory, for example as shown in a study by Garry et al., (1996) where imagining childhood events lead to participants being more likely to claim they had occurred.

In Experiments 1a and 1b, participants were forced to classify all pictures as either *old* or *new*, and had no way of indicating that the item had been encountered verbally but not pictorially. There could be at least two reasons why participants may choose the *old* response for imagined items despite having no memory of the picture itself. Firstly, they may assume that any item that was imagined had also been perceived, and make an *old* judgment on the basis of having a memory of imagining a picture. Alternatively, they may simply be incorrectly interpreting the instructions and using the *old* response to indicate that they had encountered the item in any form, as opposed to specifically having studied the picture itself. Therefore, the false memory effect in Experiments 1a and 1b may have resulted from participants knowingly choosing to use the *old* response for items that they knew they had only previously imagined.

Experiment 2 involved a source-monitoring test which gave participants the option to indicate that they had previously imagined a picture without actually having seen it. Note that unlike other studies which have used a source monitoring test in an imagination inflation procedure (e.g., Henkel & Franklin, 1998), the current procedure involves the presentation of pictures rather than word cues in the source monitoring test. Thus, as opposed to introspecting about whether a given object had been perceived and imagined or only imagined, participants study each picture and make a judgment about whether they have encountered this particular picture, or imagined a different exemplar of the same object.

The key comparison in this experiment is that between imagined and new items that are incorrectly identified as having been studied. The response options provided at test in this experiment gets around the problem of participants being forced to use the *old* response for imagined items when in fact they only remember having imagined it or seen the object's name. If participants still claim to have actually seen more of the imagined pictures than new pictures, there is stronger evidence that the false memory induction procedure has resulted in a memory of actually having studied a picture.

3.2 Experiment 2

Method

Participants

Fifty-three participants were drawn from the same pool as those in previous experiments; 29 participants were assigned to the recognition condition and 24 to the source monitoring condition.

Design and Procedure

The materials, design and procedure used in this experiment were identical to those of Experiments 1a and 1b with the exception of the test phase. The experiment was a mixed design, with item type (seen/imagined/new) as the within-subjects variable and test (recognition/source monitoring) as the between-subjects variable. The experiment was completed in one 30 minute session and involved a study phase in which participants labelled a set of pictures, an imagery phase where they tried to recall whether the pictures they had seen in the study phase had been presented in black and white or colour, and a test phase which differed between the two conditions. Apart from the test, the procedure was identical in all conditions.

Test Phase: Recognition Test. This version of the test was identical to the recognition portion of the test phase in Experiment 1a, except that there was no remember/know procedure for *old* responses. Participants were shown the 40 items one by one and made old/new judgments using the \ and / keys in their own time with no practice phase or feedback.

Test Phase: Source-monitoring Test. In this version of the test, participants were shown the 40 pictures one by one and asked to make a source judgment. Participants were instructed to indicate on each trial whether the picture was one they had *perceived and imagined* (i.e., this item had appeared as a picture in the study phase), a picture they had only *imagined* (i.e., this item had appeared in previous phases of the experiment as a word only), or a *new* picture. Keys corresponding to these options were labelled “P”, “I” and “N” respectively and responses were made by pressing the appropriate key. Participants made judgments in their own time while viewing the pictures one by one, with no practice phase or feedback.

Results

Study and Imagery Phases

Participants in the two conditions did not differ significantly on any measures of performance in the study and imagery phases. The data summarised below thus represent the performance of all 53 participants. Accuracy in the study phase was high – performance ranged from 86.0% to 100% correct identification of pictures from a choice of two labels, with a mean score of 96.5%. Participants were reasonably accurate in determining whether they had seen pictures in black and white or colour during the study phase. Responses to the 10 critical imagined items which had not appeared in the study phase were excluded from the analysis. The mean score was 62.6% (*SD* = 9.09), which was significantly different from the chance level of 50%, $t(52) = 10.11$, $p < .001$. When asked to indicate confidence on a 1 to 5 scale, participants were significantly more confident in their answers to the 80 items which had appeared in the study phase (*M* = 3.41) than to the 10 imagined items for which there was no correct response, (*M* = 2.31); $t(52) = 11.27$, *SEM* = .10, $r = .68$, $p < .001$.

Test Phase

Figure 3.1 shows the proportion of *old* responses on the recognition test to seen, imagined and new items, as well as the proportion of *perceived and imagined* responses to these items on the source monitoring test. *Old* and *perceived and imagined* responses on the recognition and source monitoring tests respectively can be seen to represent hits to seen items and false alarms to imagined and new items. An overall analysis on the *old* and *perceived and imagined* responses was conducted using a mixed ANOVA, with item type (seen/imagined/new) as the within-subjects variable and test (recognition/source monitoring) as the between-subjects variable.

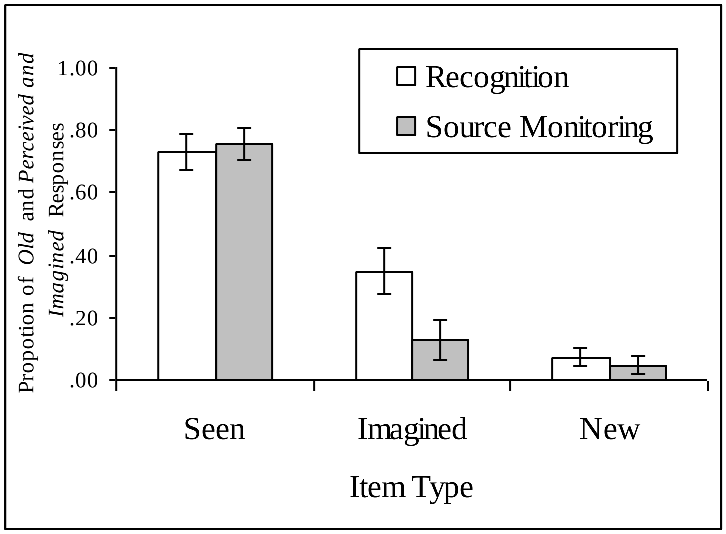

Figure 3.1: Proportion of old *responses on the recognition test and* perceived and imagined *responses on the source monitoring test to seen, imagined, and new items in Experiment 2.*

Participants showed different response patterns on the recognition and source monitoring tests: the overall ANOVA showed a significant main effect of test, $F(1, 51) = 4.65$, $MSE = .05$, *partial* $\eta^2 = .08$, $p < .05$, and an interaction between test and item type, $F(1.74, 88.78) = 8.95$, $MSE = .03$, *partial* $\eta^2 = .15$, $p < .001$ (Mauchly's test indicated that the assumption of sphericity had been violated for the main effect of the within-subjects variable, item type: $\chi^2(2) = 8.06$, $p < .05$; a Greenhouse-Geisser estimate of sphericity was used to correct degrees of freedom, $\varepsilon = .87$). In order to explore this interaction further, individual between-subjects ANOVAs were performed for seen, imagined, and new items. While there was no significant difference in the response patterns for seen or new items between the two test conditions (both *F*s < 1.5), test format had a significant effect on false alarms to imagined items, $F(1, 51) = 10.78$, $MSE = .06$, *partial* $\eta^2 = .18$, $p < .005$. Furthermore, baseline bias as defined by the criterion set in the seen-new item comparison did not differ between the source monitoring test ($M = .44$) and the recognition test ($M =$

.42; p = .39; see Appendix B). Thus, allowing participants the additional response option reduced the number of false alarms on imagined trials on the source monitoring test (*M* = .13) as compared to the recognition test (*M* = .35) without affecting the number of hits or the number of false alarms to new items.

The overall ANOVA also revealed a main effect of item type: the likelihood of participants responding *old* or *perceived and imagined* to an item was dependent upon its status, $F(1.74, 88.78) = 269.62$, $MSE = .03$, *partial* $\eta^2 = .84$, $p < .001$ (Greenhouse-Geisser, $\varepsilon = .87$). This was true for both the recognition and source-monitoring conditions, as verified by individual within-subjects ANOVAs for each condition (both *F*s > 90). Furthermore, in each condition, participants made more such false alarms to imagined items than to new items; for the recognition test, $t(28) = 5.68$, $SEM = .05$, $r = .53$, $p < .001$; and for the source monitoring test, $t(23) = 3.39$, $SEM = .02$, $r = .33$, $p < .005$. Despite having the option of assigning imagined items to a separate *imagined* response category, participants still claimed to have actually perceived these items more often than new items.

Discussion

Experiment 2 compared the standard recognition test employed in Experiments 1a and 1b with a source monitoring test in which participants were given an additional response option. Unlike in Experiments 1a and 1b where participants were forced to classify three distinct types of items into two categories, the current test included three response options which reflected the three different ways participants had interacted with the stimuli (via perception, via imagination, or not at all). When given the option to indicate that an item had only been imagined, participants still claimed to have actually perceived significantly more of the pictures that had been imagined compared to new items from the same categories. This result implies that false memory for imagined items in the current procedure cannot simply be attributed to a relaxed criterion or misunderstanding of instructions.

False memory can arise from both perceptual (i.e., data-driven, relating to surface features) and conceptual (i.e., semantic, relating to meaning) processes. For instance, in the standard DRM paradigm, the activation of critical lures is predominantly conceptual in nature, as studied words evoke the lure via semantic associations. Other studies, however, have used the same framework but achieved activation of the critical lures via perceptual processes (e.g., Sommers & Lewis, 1999; Schacter, Verfaellie, & Anes, 1997). In these experiments, participants study lists that are phonologically as opposed to semantically related to lures. In the current procedure, imagined pictures may be activated both conceptually (i.e., because the names of the objects they represent have been studied), and perceptually (i.e., because the objects are imagined). However, it is not clear which of the two makes a greater contribution to the false recognition effect. One hypothesis is that false recognition of the non-presented pictures, much like a true recognition of the pictures that had actually been perceived, arises because there is a memory trace which contains perceptual features that partially match those of that particular picture. The opposite hypothesis is that no such perceptual memory trace exists; instead, false recognition of a non-presented picture is the result of conceptual activation.

Henkel and Franklin (1998) asked a similar question in their imagination inflation study. Participants imagined objects which were either perceptually or conceptually related to other objects that they were actually shown in pictures. Perceptually related items were physically similar in form (e.g., *magnifying glass* and *lollipop*), while conceptually related items came from the same functional category (e.g., *apple* and *banana*). At test, participants were shown the names of objects which had been perceived, objects which had been imagined, and new objects, and asked to identify the source. Henkel and Franklin (Experiment 2) found that conceptual and perceptual relatedness of perceived and imagined items contributed equally to inflating incorrect claims of having perceived pictures of objects that had been imagined. Specifically, imagined objects that were similar in form to studied pictures were incorrectly attributed to prior perception as often as imagined objects that were semantically related to studied pictures, and the highest levels of such false alarms was made to pictures that were both similar in form to one picture and semantically related to another. In Experiment 3, the imagery task was adapted so that the type of processing involved was either perceptual or conceptual in nature in order to compare the relative contributions of the two processes to false recognition in the current paradigm.

3.3 Experiment 3

In the current experiment, the procedure used in Experiment 2 was modified to include two different conditions in which everything except imagery phase instructions was held constant. The imagery task in this experiment did not require participants to make judgments about how the pictures had been presented in the study phase. In the perceptual condition, participants imagined or tried to recall pictures of objects in response to their names, while in the conceptual condition, participants imagined sentences in response to the same words. Note that the perceptual condition in this case cannot be thought of as purely perceptual as it does not exclude conceptual activation arising from presentation of the word. Thus, we can only hypothesise about the potential role of perceptual processes by looking at any reduction in false recognition that arises from eliminating instructions to imagine a picture, and cannot make claims regarding whether false recognition could have arisen via perceptual processing alone. Instead, the perceptual condition could be seen as a condition which combines both conceptual and perceptual activation, which in Henkel and Franklin's (1998) study produced the highest level of source confusion. If perceptual activation contributes towards the false recognition effect observed so far, participants in the perceptual condition should claim to have perceived more of the pictures that they only studied as words than participants in the conceptual condition.

Method

Participants

Thirty-five Undergraduate students aged 18-21 (M = 19.3) from University College London took part in this study and were reimbursed for their time. Twenty-one participants were assigned to the perceptual condition and 14 participants to the conceptual condition.

Design and Procedure

The materials, design and procedure used in this experiment were based on those of Experiment 2; the differences are detailed below. The experiment was a mixed design, with item type (seen/imagined/new) as the within-subjects variable, and imagery (perceptual/conceptual) as the between-subjects variable.

Study Phase. The study phase was almost identical to those of the previous experiments, except for the number of times participants encountered the critical imagined items, which was decreased from four times in previous experiments (once in the imagery phase and three times in the study phase, as an incorrect option in the picture labelling task) to just once in the current experiment (in the imagery phase only). In order to do this, the word pool from which the incorrect options were drawn in the study phase labelling task was altered. In the current experiment, the incorrect option always corresponded to the name of another picture participants would see during the study phase, and never to the name of a critical imagined item. Although this manipulation could have the potential effect of weakening the false memory induction, the original procedure would have interfered with the imagery manipulation in the current experiment as participants are explicitly asked to imagine pictures in response to the verbal labels in the study phase task.

Imagery Phase. The general format of the imagery phase, including presentation times, was similar to those of previous experiments: participants were presented with words for 4000 msec each and made a response to each word. As in previous experiments, participants were informed that the words they were about to see related to the pictures they had seen earlier on. In the perceptual condition, participants were asked to imagine the picture that the word referred to, and rate the pleasantness of this image. In the conceptual condition, participants were asked to think of a sentence using the word, and rate the pleasantness of the sentence. Aside from these instructions, there was no difference between conditions. In both conditions, participants made responses on the same scale, from 1 = *not at all pleasant* to 5 = *extremely pleasant*.

Test Phase. This phase was the same in the two conditions, and test items were identical to those in previous experiments. Participants were shown the 40 pictures one by one and asked to make a source judgment. They were instructed to indicate whether they had studied the picture AND its label, its label only, or whether it was a new item that they had not previously encountered in the experiment. The response options were the equivalent of those in Experiment 2, so *perceived and imagined* became *picture and label* while *imagined* became *label*; this was done so as to make the response options appropriate to the conceptual condition, where participants had not imagined a sentence instead of a picture. Keys corresponding to these options were labelled “P”, “L” and “N” respectively and responses were made by pressing the appropriate key. Participants made judgments in their own time while viewing the picture, with no practice phase or feedback.

Results

Study and Imagery Phases

Participants in the two conditions did not differ significantly on any measures of performance in the study and imagery phases. The data summarised below thus represent the performance of all 35 participants. Accuracy in the study phase was high as in the previous experiments, with a mean score of 97.0% across the two imagery conditions; there were no differences in this measure between conditions ($p > .6$). Pleasantness ratings in the imagery task were almost identical across the two imagery conditions ($M = 3.18$, $p > .9$) and did not differ depending on whether participants had seen pictures referring to the words in the study phase (i.e., between the 10 imagined items and the 80 items that had appeared in the study phase; $p > .5$).

Test Phase

Figure 3.2 shows the proportion of *picture and label* responses to items in the source monitoring test by imagery condition.

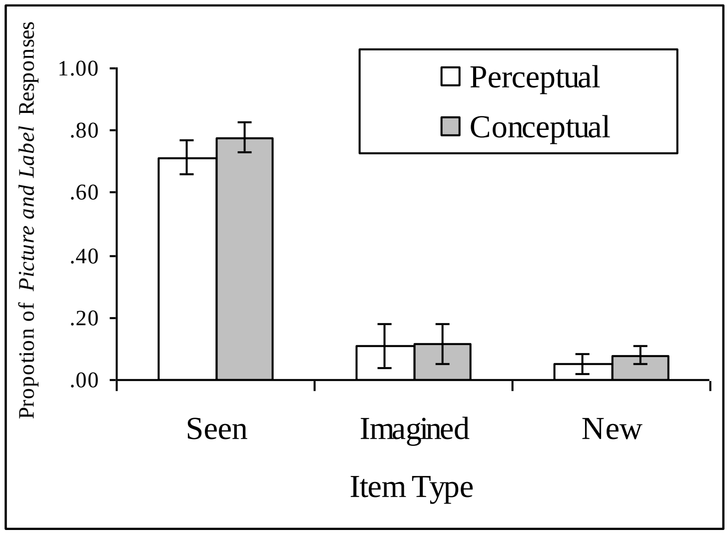

Figure 3.2: Proportion of picture and label *responses to seen, imagined, and new items in the perceptual and conceptual imagery conditions in Experiment 3.*

An overall analysis on these responses was conducted using a mixed ANOVA, with item type (seen/imagined/new) as the within-subjects variable and imagery (perceptual/conceptual) as the between-subjects variable. Across the two conditions, the likelihood of participants responding *picture and label* to an item was dependent upon its status, $F(1.43, 47.33) = 269.38$, $MSE = .03$, *partial* $\eta^2 = .89$, $p < .001$ (Mauchly's test indicated that the assumption of sphericity had been violated for the main effect of the within-subjects variable, item type: $\chi^2(2) = 16.06$, $p < .001$; a Greenhouse-Geisser estimate of sphericity was used to correct degrees of freedom, $\varepsilon = .72$). The proportion of hits to pictures which had appeared in the study phase ($M = .74$) as well as false alarms to pictures which had appeared as words in the imagery phase ($M = .11$) were both significantly higher than the proportion of false alarms to new pictures ($M = .06$; $t(34) = 18.46$, $SEM = .04$, $r = .91$, $p < .001$ and $t(34) = 2.51$, $SEM = .02$, $r = .19$, $p < .05$ respectively).

Crucially, the imagery manipulation did not have an effect on responses and there was no interaction between imagery and item type (*F*s < 1). In other words, whether participants imagined pictures (perceptual) or sentences (conceptual) in the imagery phase did not affect the proportion of false alarms they made to these items on the picture source monitoring test. The bias data for the seen-new comparison indicate that although the response criterion was numerically lower in the conceptual condition (M = .27) than the perceptual condition (M = .40), this difference did not reach significance (p = .33; see Appendix B). A cross-experimental comparison confirmed that the number of times items that had only appeared as words were incorrectly deemed to have been perceived did not differ significantly between Experiment 2 where participants encountered the words four times (M = .13), and the current experiments where participants encountered them only once in the imagery phase (M = .11; p = .65). An examination of the *imagined* (Experiment 2, M = .75) and *label* (current experiment, M = .61) responses to these items, on the other hand, indicates that participants benefited significantly from having encountered the words in the study phase when it came to correctly identifying them at test, $t(57) = 2.41$, SEM = .06, *partial* η^2 = .30, $p < .05$.

Discussion

Experiment 3 fulfilled two aims: replicating the pattern of results found in Experiment 2 with two new imagery tasks, and evaluating the contribution of perceptual and conceptual processes to the false recognition effect. In the source monitoring test, participants presumably attempt to match the attributes of the stimulus to their memory of the item, and will indicate that they had previously perceived the picture if there is a match between the two. There are thus two possible factors that might lead to participants making false alarms to imagined items. One factor, present in both the conceptual and perceptual versions of the imagery task, is the conceptual activation of the object depicted in the picture that arises from having studied its name. Another factor, which occurs in the perceptual condition only, is perceptual activation arising from imagination of the object. If the image participants create in response to a particular word is sufficiently similar to the picture itself to produce a strong match between the stimulus and the memory of the imagined item, this might lead participants to believe they had perceived the picture. If perceptual activation was a significant driver of the false recognition effect, we would have expected to see higher levels of false recognition in the perceptual as compared to the conceptual imagery condition. However, participants made no fewer *picture and label* responses to imagined items when they were instructed to imagine sentences instead of pictures, thus eliminating the perceptual component of the task.

Two caveats are to be noted here. Firstly, the conceptual imagery task could have been contaminated by perceptual processes. Although there were no instructions in the conceptual condition relating to imagination of pictorial representations, participants may have spontaneously imagined pictures of the object while reading the object's name and imagining a sentence. There is indeed some evidence that participants spontaneously produce images of objects as a result of interacting with the objects' names (e.g., Foley, Durso, Wilder, & Friedman, 1991; Foley & Foy, 2008). Secondly, the imagery task in Experiment 3 was manipulated with reference to Henkel and Franklin's (1998) perceptual/conceptual similarity manipulation, where the additive condition in which the imagined object was similar to a studied picture both semantically and in terms of perceptual features produced the highest level of source confusion. However, there was a key difference between their procedure and the one employed here. Henkel and Franklin investigated the contribution of perceptual and conceptual processing without manipulating the imagery task participants engaged in, while in Experiment 3 type of processing was manipulated by giving participants different imagery instructions: in the conceptual condition, participants were instructed to imagine a sentence, while in the perceptual case, participants were instructed to imagine a picture. Intraub and Hoffman (1992) found that reading a paragraph describing a visual scene lead to more incorrect claims of having studied the associated scene itself when there were no imagery instructions. According to Intraub and Hoffman's interpretation of the source monitoring framework, the explicit picture imagination instructions in the perceptual condition decrease source confusion between seen and imagined pictures because the memory of cognitive operations involved in the imagination process serve to alert participants to the fact that they had generated the image rather than seen it. In the case of spontaneous imagery, on the other hand, there will be no record of such cognitive processes and thus no cue that the image was self-generated. The use of such a monitoring strategy could explain why there was no difference in false recognition between the conceptual and perceptual conditions.

3.4 Discussion of Experiments 2 and 3

In Experiments 1a and 1b, participants may have made false alarms to imagined items because they misinterpreted the *old* response option on the recognition test and called pictures *old* if they remembered having studied the associated word. The results of Experiments 2 and 3 discredit the hypothesis that participants are always aware of not having studied the pictures. In other words, there is evidence to suggest that in at least some cases, when shown a picture of an object that had only appeared as a word, participants will actually recognise this picture as having been studied. However, at least some of the effect is likely to be driven by diffuse familiarity, as evidenced by the difference in false alarms to imagined items in Experiments 2 and 3 compared with Experiments 1a and 1b.

Experiment 3 demonstrated that the false recognition effect was not dependent on the specific imagery task employed in Experiments 1a and 1b. Even when asked to perform a conceptual task on the object name (imagining a sentence which used that word), participants were just as likely to incorrectly claim they had seen the associated picture. The finding that imagining a picture is no more effective in making people think they had studied it than imagining a sentence leads to the tentative conclusion that perceptual processes are not responsible for this effect. However, it does not rule out the possibility that false memories of non-presented items do contain perceptual features. This possibility was further tested in Chapter 6. In Chapter 4 (Experiment 4), another type of direct test was used: participants were presented with two pictures side by side and asked to indicate which of the two pictures had been studied. The key question is whether participants will respond to studied pictures differently when these are paired with either imagined or new pictures.

CHAPTER 4: FALSE MEMORY FOR PICTURES ON A FORCED CHOICE TEST

So far, four experiments have successfully demonstrated a new procedure that rapidly induces false memory for distinctive pictures. In each experiment, attempts have been made to verify that this effect is not solely due to a relaxed response criterion. In Experiment 1a, participants not only claimed to have seen imagined items, but also indicated that they remembered their specific occurrence in the study phase. In Experiment 1b, participants endorsed imagined items as old with high confidence. Finally, in Experiments 2 and 3, participants indicated that they had actually seen imagined pictures even when they could have classified them as imagined. These data all point towards the conclusion that participants are reporting a genuine false memory. However, it is still possible that participants claimed to have seen pictures of imagined objects when they only remembered having studied their names because of the high correlation between these two events. In Experiment 4, an even stricter test was used to reject this possibility.

In a forced choice test, participants are presented with pairs of items, usually where one is old and the other new, and asked to indicate which of the two they had studied. This test is considered to be superior to standard old/new recognition because it is not affected by the criterion set by participants (Macmillan & Creelman, 2005; Shepard & Chang, 1963). In addition, this task is also much simpler for the participant than a test in which each item has to be classified individually. In fact, this is the task used by Shepard (1967) in his famous study where participants achieved 99.7% correct recognition of pictures after a 2-hour delay.

Consider the experiments presented here so far: on seen trials, participants may correctly and with great ease identify the pictures they recall seeing, but on imagined trials where they recall seeing a word but not the associated picture, they may apply a rule that allows an old response in cases where the memory of the associated word is particularly strong; the two strategies are not mutually exclusive. In a forced choice test, this sort of dual-basis decision strategy is impossible. Therefore, unless there are genuine false memories of imagined pictures, participants should choose the seen picture as often on trials where a seen picture is presented alongside an imagined picture, as on trials in which it is presented alongside a new picture.

4.1 Experiment 4

The present experiment involved a two-alternative forced choice (2AFC) test which allowed pairs of items to be compared directly. Participants were presented with two items side by side and asked which of the two pictures they had previously studied. Each combination of items (seen-imagined, seen-new, and imagined-new) was tested; in seen-imagined and seen-new pairs, only the seen item had been studied and was therefore the correct response, while in imagined-new pairs, neither item had been studied so there was no correct response. The prediction for imagined-new pairs does not depend on whether participants genuinely remember having studied the imagined pictures or only remember having studied the associated word: in either case, participants should choose the imagined item more often. While the memory explanation attributes this choice to an underlying memory of the imagined item, it is also reasonable to speculate that in the absence of a memory for either picture, participants will either guess at random or choose the imagined item due to having been exposed to it in some form.

In the absence of memory for imagined pictures, seen-imagined and seen-new pairs should yield the same level of accuracy. Based on the results from previous experiments, however, it is possible that participants may find it more difficult to correctly select the studied item from seen-imagined than seen-new pairs, even though in both cases there is an equal chance that they will have correctly identified the picture as having been studied if it had been presented alone. This would be a highly counterintuitive result, and would provide very strong evidence that the false alarms to imagined items obtained in the current procedure reflect genuine memory illusions.

Method

Participants

Thirty-nine undergraduates aged 18-20 (*M* = 18.7) participated in the experiment as part of a course requirement.

Materials

The test phase in the current experiment required double the number of critical items used in the previous experiments reported here. For this reason, the number of items in each phase of the experiment was modified, although the same corpus of 110 categorised pictures was used. The number of critical items was increased from 30 to 60, and these 60 pictures were divided into three groups of 20 items which were rotated through the seen, imagined, and new conditions. The remaining 50 non-critical items appeared as pictures in the study phase only for participants in every condition. In the study phase, participants were presented with a list of 70 pictures, consisting of the 20 seen critical items and the 50 non-critical items. In the imagery phase, participants studied 70 words: 20 seen items, 20 imagined items which replaced 20 of the non-critical items from the study phase, and 30 filler items from the study phase. Finally, in the test phase, participants all saw the same 60 critical items, which were presented in pairs of seen-new, seen-imagined, and imagined-new items (see below).

Design

Pair type (seen-new/seen-imagined/imagined-new) was manipulated within subjects, and the dependent measure was the proportion of trials on which participants chose the first item in each pair for each pair type.

Procedure

The present experiment was similar to Experiments 1-3 in terms of structure (study phase, imagery phase, test phase), with the modifications detailed below.

Study Phase. Aside from the items (see *Materials*), the study phase was identical to that of Experiment 3. The incorrect option in the labelling task always corresponded to another picture participants would see during the study phase, and never to one of the imagined items. Participants thus did not encounter the imagined items until the imagery phase. The imagery phase instructions and procedure were identical to those of Experiments 1-2, where participants were shown words and asked to recall whether the associated picture had appeared in colour or in black and white in the study phase.

Test Phase. The test phase involved making judgments on pairs of pictures. Participants were shown two pictures side by side and asked to indicate which of the two pictures (left or right) they had previously studied. Following this judgment, participants indicated their confidence on a scale of 1 to 5 where 1 = *not at all confident* and 5 =

extremely confident. There were 10 each of three types of trials: seen-new, seen-imagined, and imagined-new, and the position (left or right) of the items was randomised with the constraint that each item within each pair appeared in each position exactly five times.

Results

Study and Imagery Phases

Performance in the study phase ranged from 91.4% to 100% correct identification of pictures from a choice of two labels, with a mean score of 97.1%. Mean performance in the imagery phase, where participants had to indicate if the picture had appeared in colour or in black and white in the study phase, was 63.6% ($SD = 8.9$) for the 50 items which had appeared in the study phase. This was roughly equivalent to performance in the imagery phase in Experiments 1 and 2, and significantly different from the chance level of 50%, $t(38) = 9.62$, $p < .001$. On the 1 to 5 confidence scale, participants were significantly more confident on their answers to seen items ($M = 3.55$) than to imagined items ($M = 1.94$); $t(38) = 12.04$, $SEM = .14$, $r = .68$, $p < .001$.

Test Phase

All three stimulus pairs (seen-new, seen-imagined, and imagined-new) were used in the analysis. Figure 4.1 shows the probability of choosing the first item in each pair for each pair type, which represents accuracy in the case of the seen-new and seen-imagined pairs. The probability of choosing the correct option from a seen-imagined pair ($M = .86$) was lower than that of choosing the correct option from a seen-new pair ($M = .96$; $t(38) = 3.70$, $SEM = .03$, $r = .40$, $p < .001$, indicating that participants found it more difficult to correctly identify the studied picture when they had previously imagined the other picture in the pair than when the other picture was new. For the imagined-new pair, a one-sample t-test confirmed that participants chose the imagined picture at .22 above the chance level of 0.50, $t(38) = 7.94$, $SEM = .03$, $p < .001$.

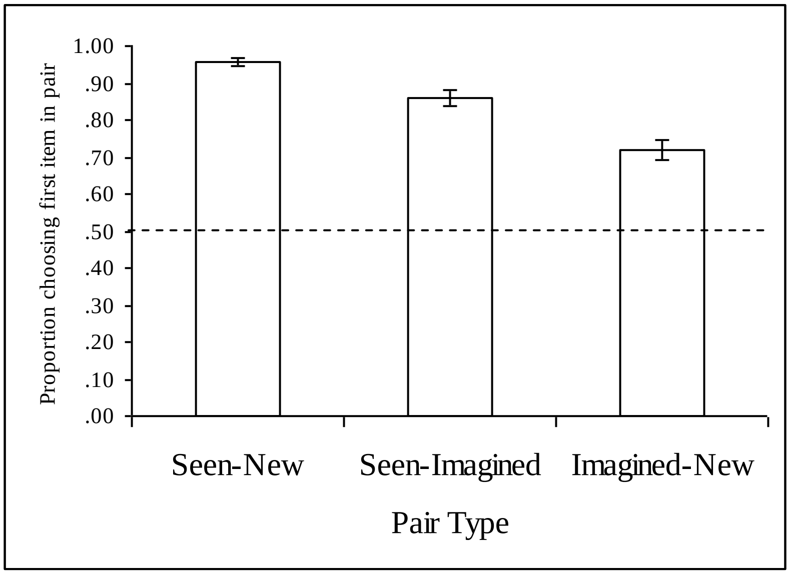

Figure 4.1: Probability of choosing the first item in each type for seen-new, seen-imagined, and imagined-new pairs in Experiment 4. The dashed line represents chance performance.

Due to the low number of incorrect trials, confidence was compared by pair type rather than by response. Participants did not differ significantly in the confidence ratings they assigned to seen-new pair decisions ($M = 4.01$) and seen-imagined pair decisions, ($M = 3.91$, $p = .28$), but were far less confident on imagined-new pairs (M = 2.27; compared with seen-new pairs, $t(38) = 10.45$, $SEM = .17$, $r = .75$, $p < .001$). The time taken for participants to make a response on each trial was also recorded and is shown broken down by pair type in Figure 4.2 (data were removed for the two participants whose mean time to respond across all three pair types was more than two standard deviations away from the mean). Participants deliberated on average 269 msec longer on seen-imagined trials than on seen-new trials, $t(36) = 2.03$, $SEM = 132$, $r = .16$, $p < .05$.

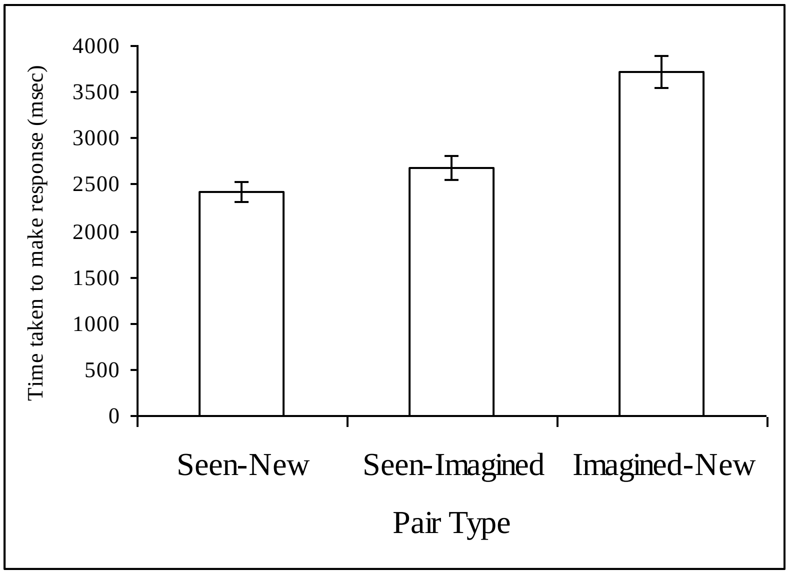

Figure 4.2: Time taken to make a decision on seen-new, seen-imagined, and imagined-new pairs in Experiment 4. These data are based on 36 of the 38 participants due to outliers.

Discussion

In Experiments 1-3, participants evaluated items individually and consistently responded differently to imagined and new items. In Experiment 4, comparisons were made between two items presented side by side. Of most interest is the comparison between seen-imagined and seen-new pairs: participants performed seen-imagined decisions with less accuracy than seen-new decisions. With respect to the seen picture, the task of selecting this picture from a seen-imagined pair should be equivalent to that of selecting it from a seen-new pair. The likelihood that the seen picture was encoded and would have been recognised on a standard old/new recognition test is equal between the two conditions. In the current procedure, on the other hand, one factor which can make a forced-choice decision more difficult is the similarity between alternatives. For instance, in the DRM paradigm, one might predict that deciding between a list word and an unrelated item will always be easier than deciding between a list word and the critical item, simply because of the semantic relationship between the two (there have, in fact, been no published studies applying the forced choice test procedure to the DRM paradigm, but see McDermott, Chan, & Weinstein (2009) for some initial findings). In the current procedure, on the other hand, the items were fully counterbalanced, so the two items in a seen-imagined pair would always be equally similar in terms of perceptual features compared to those in a seen-new pair. Nevertheless, participants found it harder to correctly select the seen picture when the name of the object in the alternative picture had been presented in the imagery phase than when the alternative object have never been named. Note that the accuracy rate for seen-new pairs was extremely high, and very similar to that obtained by Shepard (1967) in his much-cited study showing ceiling levels of recognition performance for pictures in a forced choice procedure.

An argument can be made for predicting lower accuracy on seen-imagined trials compared to seen-new trials that is not contingent on participants having a genuine false memory of the imagined pictures. This argument is based on the assumption that participants did not encode some of the pictures in the study phase (e.g., because they were not paying attention at the time they were presented), and have to make a guess on any trials where these appear alongside another picture that they do not recognise. In the case of seen-new trials, the two pictures should appear equally familiar and participants will thus select each of the two items with an equal probability of .5. In the case of seen-imagined trials, on the other hand, the imagined item may feel more familiar because its name was encountered in the imagery phase, so participants may default to choosing this item.

Applying this logic to the current data, given the .04 error rate for seen-new pairs, participants may actually have been guessing on .08 of seen-new trials (i.e., on twice as many trials than the number they got wrong, because a guess yields a .5 success rate), which, when rounded up to the nearest number of trials, amounts to 1 trial per participant. If the same number of guesses were made on seen-imagined trials, and participants always chose the imagined item in these cases[2], we would expect the accuracy rate for seen-imagined trials to be around .9 (i.e., 9 out of 10 trials); in fact, it was significantly lower than that ($t(38) = 1.77$, $SEM = .02$, $p < .05$, one-tailed). Even with this strict criterion, there is evidence that studying the name of a picture and imagining it interfered with participants' ability to reject it in favour of another picture which had actually been seen 15 minutes previously.

CHAPTER 5: FALSE MEMORY ON INDIRECT TESTS: A LITERATURE REVIEW

While memory can be investigated via both direct (or explicit, conscious, intentional) and indirect (or implicit, unconscious, unintentional) tests, most of the research into false memory so far has employed direct recognition or recall tests. In all the procedures described here up to this point, participants explicitly report having seen items they did not study. Indirect tests, on the other hand, do not depend on participants using a conscious decision strategy. Instead, prior exposure facilitates performance on certain tasks including lexical decision (Ratcliff et al., 1985) and picture naming. Recently, researchers have investigated whether false memory for critical lures in the DRM paradigm is exhibited on indirect as well as direct tests. Given the limited amount of research on the subject, this chapter is devoted to a thorough review of all the studies which have addressed this issue.

The striking result in the standard version of the DRM paradigm is that true and false memories are reported at equal rates and thus appear to be indistinguishable by a simple recognition test. While it has already been shown that true and false memories tested directly can be dissociated using procedures such as repeated study (Robinson & Roediger, 1997), it is interesting to see whether indirect tests can also help distinguish between the two. The most intuitive hypothesis is that non-presented critical lures should not facilitate performance on an indirect test that benefits from prior perception. However, if false memories can acquire perceptual details (e.g., Lampinen et al., 2005), a contrasting hypothesis would be that activation of related items could lead to priming by a similar mechanism to perception. In Chapter 6, the new procedure proposed in this dissertation is used to test whether the false memory of an imagined picture behaves like a true memory of a perceived picture on a perceptual indirect test.

Indirect tests can be roughly divided into those that tap into the perceptual and conceptual details of studied stimuli. Indirect perceptual tests such as speeded recognition of degraded stimuli rely on the identification of perceptual features, while indirect conceptual tests such as word association access an item's meaning. It has generally been found that performance on indirect conceptual tests is facilitated for critical lures as well as for studied items, while the behaviour of critical lures on indirect perceptual tests is more controversial. All the studies so far that have investigated this question (often referring to the phenomenon as "false implicit memory") have used verbal rather than pictorial stimuli. Some studies claim to have found significant priming on perceptual tests for critical lures even though they were never perceived at study (e.g., McKone & Murphy, 2000), but other studies have challenged such results by conducting experiments with stricter controls for awareness (e.g., McBride, Coane, & Raulerson, 2006). Convincing results showing perceptual priming of critical lures equivalent to that of studied words could lead to the strong conclusion that true and false memories in the DRM paradigm do not differ in terms of perceptual attributes. The main findings of this line of enquiry are discussed below, followed by an empirical extension in Chapter 6 to pictorial stimuli using the procedure presented in previous chapters.

5.1 False Memory for Verbal Stimuli on Indirect Tests

False memory for verbal stimuli has been investigated using the following indirect tests: word-stem completion (Hicks & Starns, 2005; McBride et al., 2006; McDermott, 1997; McKone & Murphy, 2000); word-fragment completion (McDermott, 1997; McKone & Murphy, 2000); lexical decision (Hancock, Hicks, Marsh, & Ritschel, 2003; McKone, 2004; Meade, Watson, Balota, & Roediger, 2006; Tse & Neely, 2005; Whittlesea, 2002; Zeelenberg & Pecher, 2002; word association (McDermott, 1997); word naming (Whittlesea, 2002 perceptual identification (Hicks & Starns, 2005); and anagram solving (Lövdén & Johansson, 2003).

McDermott (1997) found that priming of critical lures on a conceptual test (word association) was equivalent to priming of studied items, while priming for critical lures on word-stem completion and word fragment completion tests (both putative perceptual tests) was significant but generally lower than that of studied items. In the word-stem completion test, subjects are presented with three-letter stems such as *CHA__* and, without reference to the study phase, instructed to produce the first suitable word that comes to mind. Memory is then assessed by the number of times a stem is completed by a particular word under conditions where this word had or had not appeared in the study phase (or, in the case of a critical item in the DRM paradigm, where its associates had or had not appeared). The word fragment completion test is scored in the same way, but test probes are partially completed words such as *C_AI_*. McKone and Murphy (2000) replicated word-stem and word fragment completion findings with stricter methodological controls. These authors used word norms to match critical lures and list items, which is not done in the standard DRM procedure, and excluded participants who claimed to use explicit retrieval strategies. However, McBride et al. (2006) found significant priming for lures on word-stem completion only in those participants who were deemed "test-aware" (20/54 participants) because they had noticed that some of the word-stems could be completed with items from the study phase.

The transfer-appropriate processing (TAP) framework proposes that performance on indirect tests that rely on perceptual processes should be sensitive to changes in presentation modality. For instance, Blaxton (1989) showed that priming on a word-fragment completion task was greatest when presentation modality (auditory/visual) was kept constant between study and test. Is priming of critical lures on the word-stem completion task modality-specific? If critical lures are primed in the same way as studied items, they should also be affected by a modality change from study to test. For instance, critical lures of lists studied aurally should be less facilitated on a visual word-stem completion test than critical lures of lists studied visually. If this is not found to be the case, the priming exhibited for critical items on perceptual indirect tests must arise from non-perceptual activation.

McKone and Murphy (2000) compared priming of critical lures in participants who had studied either visual or auditory lists. In this between-subjects design, priming of critical lures on a visual word-stem completion task was reduced by a mismatch between study and test modality. Hicks and Starns (2005), on the other hand, found that while critical lure priming on a visual word-stem-completion task was significant (though lower than list item priming), this priming was unaffected by a within-subjects manipulation of modality (visual/auditory), unlike priming of list items. Hicks and Starns suggested that if priming of critical lures can withstand a change in modality, the representations of these unstudied items could be non-specific in terms of modality and thus unlike the representations of studied items. Instead of being based on perceptually encoded information, critical lures could be activated conceptually. Although word-stem completion is considered to be a perceptual test, it is more conceptually driven than perceptual identification (Marsolek & Andresen, 2005). Using a version of the latter, where words were presented for identification for 24ms followed by a backward mask, Hicks and Starns found priming for critical lures only when these were actually included in the study lists and not when they were activated using the standard DRM procedure. Lövdén and Johansson (2003) also found evidence to suggest that critical lure priming in so-called perceptual tasks is not dependent on visual representations. They found significant facilitation for critical lures on an anagram task, but this priming was eradicated when articulation was suppressed. Whether priming of critical lures is modality-specific addresses the same issue as studies that examined whether false memories tend to be reported in the same modality as list items (e.g., Hicks & Hancock, 2002). The conflicting findings reviewed above suggest that while the representations of false memories may not be as perceptually rich as those of true memories, it is possible that some activation of perceptual attributes does occur, leading to facilitated performance on indirect perceptual tests.

Lexical decision is another indirect test that is generally considered to be perceptually based and has produced mixed results for critical lures in the DRM paradigm. In lexical decision, subjects must classify probes as words or non-words at speed. There has been some inconsistency in the findings as to whether performance on lexical decision is facilitated for non-presented critical lures. Zeelenberg and Pecher (2002) found no priming of critical lures on a lexical decision task using McDermott’s (1997) procedure which resulted in significant priming on two other perceptually-based indirect memory tasks described above, word-stem and word-fragment completion. McKone (2004) also found no priming of critical lures on lexical decision using the normalised stimuli from McKone and Murphy (2000).

The studies described above administered the lexical decision task at the end of the entire study period, once all the lists had been studied, and found no priming for non-presented critical lures. Three other studies administered separate lexical decision tests after study of each list. This procedure is designed to reveal the automatic spreading activation that is thought to be responsible for the formation of false memories. Hancock et al. (2003) found facilitation for critical lures on a lexical decision task compared to control words matched for length and word frequency. Subjects performed lexical decision even quicker for lures than for list items when 14 associates per list were studied, and at the same speed as list items on shorter lists. Tse and Neely (2005) replicated this effect using more rigorous counterbalancing in order to produce a more appropriate baseline reaction time (using the Balota et al. (2002) lexicon). Finally, Meade et al. (2006) did not find priming of critical lures on a more strict lexical decision task administered immediately after each study list except when the lure appeared as the first item on the test.

Whittlesea (2002) reported the only study to have found significant priming of critical lures on lexical decision following study of multiple DRM lists. Whittlesea found significant critical lure priming on a lexical decision task but no facilitation on a speeded naming task for the same critical lures. In this case, the lexical decision task was actually presented as a conceptual task, when compared with speeded naming which was considered to be more perceptual in nature. The result was attributed to the fact that lexical decision employs relatively more conceptual processes than speeded naming. Whittlsea used the *discrepancy-attribution hypothesis* to explain the false memory effect as a misattribution of the surprising clash between fluency in one domain (conceptual) and lack of fluency in another (perceptual) to prior perception. However, Whittlesea used an unconventional procedure whereby subjects studied all associates on a list together rather than being shown each word individually, so these results cannot be taken as conclusive evidence of critical lure priming in lexical decision. So far, evidence seems to suggest that the mechanism behind false recognition in the DRM paradigm does not guarantee facilitation on a word identification task.

The studies described above differ in terms of testing procedures and controls for explicit contamination. Another detail which has not been kept constant is the encoding instruction (intentional/incidental). Hicks and Starns (2005) did systematically manipulate encoding instructions and found no effect of intentional versus incidental encoding on word-stem completion priming, which was significant for critical lures in both cases. On the other hand, McBride et al. (2006) were the only authors to have conducted all their experiments with incidental study instructions, and they did not find significant priming of critical lures on word-stem completion in test-unaware subjects. It may be that intentional encoding instructions could have increased the likelihood of explicit contamination in some of the studies cited above and thus inflated any critical lure priming. In sum, the question of whether critical lures in the DRM paradigm behave like old or new items on indirect tests has not been definitively resolved. Note that the above has been a predominantly empirical review, because the theories typically used to explain the false memory effect on direct tests (e.g., Miller & Wolford, 1999; Roediger et al, 2001) do not make specific predictions as to whether a falsely remembered item would generate facilitated performance on an indirect test.

5.2 False Memory for Pictorial Stimuli on Indirect Tests

The studies reviewed above have all used indirect tests to investigate false memories for verbal stimuli. As yet, no analogous experiments have been carried out using pictorial stimuli. The aim of the experiments presented in Chapter 6 was to extend the research into false memory on indirect tests to pictures. In the current procedure, false recognition of imagined items may occur because these items are processed more fluently than new items, leading participants to misidentify them as studied items. Support for this idea comes from studies that manipulated processing fluency at test and found that this fluency was misattributed to prior perception (Jacoby & Whitehouse, 1989). Perceptual priming is one behavioural manifestation of such fluency. If perceptual priming is found alongside false recognition in the current procedure, this would provide another example of perceptual fluency incorrectly attributed to previous perception of an item. The line of enquiry in Chapter 6 also resumes that of Experiment 3, which involved a comparison of the contributions of perceptual and conceptual imagery processes to the false recognition effect. In that experiment, there was no evidence that perceptual processes made a significant contribution to the effect. A finding of no perceptual priming of imagined items would support this result and strongly suggest that the false recognition effect reported here is conceptual in nature.

CHAPTER 6: FALSE MEMORY FOR PICTURES ON AN INDIRECT TEST

It is typically found that pictures that have been previously perceived are recognised faster than new pictures (Warren & Morton, 1982), so it is expected that in the current procedure seen pictures will be recognised significantly faster than new pictures. Of greater interest is performance on pictures which had been imagined but not actually perceived. The predictions of the TAP framework mentioned in Chapter 5 (Blaxton, 1989) in relation to modality change in word priming also apply to a change in form. If priming is a function of the match between study and test, then a change in form (in this case, from verbal to pictorial) should reduce priming on perceptual implicit tests. While most studies have found that a change from verbal stimuli in the study phase to pictorial stimuli in the test phase eliminates perceptual priming completely (Hirshman, Snodgrass, Mindes, & Feenan, 1990; Scarborough, Gerard, & Cortese, 1979; Warren & Morton, 1982; Weldon et al., 1995), some other studies (Durso & Johnson, 1979; Park & Gabrieli, 1995; Srinivas, 1993) have found significant cross-form priming, suggesting that certain ostensibly perceptual priming tasks can include a nonperceptual component. In order to test whether the imagination manipulation used in the current procedure creates false memories that are identifiable on an indirect perceptual priming test, it was first necessary to select a task which was, as far as possible, data-driven and unsusceptible to nonperceptual processes.

6.1 Experiment 5

The selected indirect task was a version of the continuous identification procedure (Feustel, Shiffrin, & Salasoo, 1983; Stark & McClelland, 2000) that has been shown to be highly data-driven when used with verbal stimuli. In Experiment 5, a simplified procedure was used in which participants studied a mixed list of pictures and words, and were later tested on their memory for pictures both indirectly (with the continuous identification procedure) and directly (with the same recognition task as in Experiment 1b). This experiment sought to verify that priming on the continuous identification task for pictures only results from studying pictures themselves, and not studying their names in the absence of imagery instructions.

Method

Participants

Twenty Undergraduate students aged 18-21 (M = 19.6) from University College London volunteered to participate in the experiment in return for course credit.

Materials

The 110 images used in previous experiments were divided into three groups of 30 critical items which were rotated through the presentation format conditions (seen/word/new), and 20 filler items which appeared in the study phase only. Three counterbalancing conditions were required for this. Each 80-item study list consisted of 30 words and 30 pictures which later appeared in the test, and a further 10 filler items in each presentation format. The 90-item test phase consisted of 30 items which had been studied as pictures, 30 items which had been studied as words, and 30 new items. Half the items appeared in colour and half in black and white in the study and test phases; this factor did not play a role in the current experiment but was retained for consistency with the rest of the experiments. For the priming test, two versions of a mask were created by randomly filling a 300 x 300 pixel grid with 36 fragments taken from pictures not shown at test (see Figure 6.1).

Figure 6.1: Masks used in the continuous identification task. The two masks were randomly alternated between trials and appeared in black and white or colour depending on the colour format of the stimulus.

Design

The experiment was a repeated measures design with item type (seen/word/new) as a within-subjects variable, and two dependent measures, proportion judged as *old* in the recognition test and reaction time (RT) to identify pictures in the priming test.

Procedure

This experiment consisted of two phases: a study phase in which participants were presented with pictures and words, and a test phase in which participants performed perceptual identification followed by recognition of pictures on each trial.

Study Phase. In the study phase, participants were presented with a list of 40 pictures and 40 words. Participants were not given any particular encoding instructions; they were simply asked to watch the picture and word sequence carefully. Picture and word trials were interspersed randomly. Picture trials consisted of (a) a blank interval of 1000 msec, (b) a fixation cross appearing at the centre of the screen for 1000 msec, (c) the picture presented at the centre of the screen for 250 msec, and (d) a blank interval of 1000 msec. Word trials consisted of a blank interval of 1000 msec followed by the word presented at the centre of the screen for 4000 msec. These timings were used in order to keep presentation times of the two item types consistent with those used in all other experiments.

Test Phase. Following a 5-minute mental arithmetic distractor task, participants were given instructions for the test phase. On each trial, participants first identified a picture using the continuous identification procedure described below. Following its identification, the picture was displayed again and participants made a recognition judgment in their own time. Pictures were presented in the same central location as during the study phase. Previously encountered pictures appeared in the same colour format as in the study phase, while new pictures appeared in black and white or in colour according to counterbalancing condition, so that half of the items seen at test were in black and white and half were in colour.

The picture identification sequence used in the test phase was based on the CID (continuous identification) procedure employed by Stark and McClelland (2000) and originally formulated by Feustel et al. (1983) for use with words. In the present experiment the procedure was redesigned for pictures. The instructions stated that this part of the experiment was concerned with visual perception and that we were interested in the time it took people to accurately identify pictures. Participants were told that on each trial they would see a picture presented on the screen for a very brief duration, followed immediately by a mask. They were told that the picture would then reappear and continue to flash on the screen for longer and longer durations, making it appear clearer over time. The task was to try to identify the picture as quickly as possible. As soon as participants could make out the picture they were to press ENTER and type in the name of the picture when prompted. They were informed that once they pressed ENTER the mask would reappear and they would not see that picture again until they had entered a response. Participants were told to answer as fast as possible but only to do so once they were confident that they would identify the picture correctly. As soon as the ENTER key was pressed, RT measurement stopped, and once ENTER was pressed again to complete the response, the recognition sequence of the trial initiated.

The timing of the procedure was programmed using the ExacTicks code for Visual Basic. Each trial began with the presentation of one of the two masks (see Figure 6.1) for 500 msec. The two masks were rotated randomly between stimuli, and always appeared in the same colour format as the corresponding stimulus. Following initial presentation of the mask, the stimulus was presented for 17 msec (the duration of one screen refresh), and the mask followed for 233 msec, making a 250 msec block. This stimulus presentation was repeated with the duration of the stimulus increasing by 17 msec on every presentation, while total block time remained constant (see Figure 6.2 for a graphic representation of this procedure).

Upon identification of a picture, participants were shown the picture again and asked to judge whether it was *old* or *new*. This task was identical to the recognition test in Experiment 1b: recognition judgments were made on a confidence scale from 1 to 6 where 1 = *very sure new* and 6 = *very sure old*. Participants were specifically instructed to only use the *old* response if they had seen the picture itself in the study phase. Unlike in all the other experiments presented here, instructions stated that if participants remembered seeing the word in the study phase, the correct response to the picture was *new*. Participants made the recognition judgment in their own time while viewing the picture. Before beginning the test, participants performed a practice block of seven identification trials, although no recognition judgments were required. No feedback was given to participants about their performance during the test.

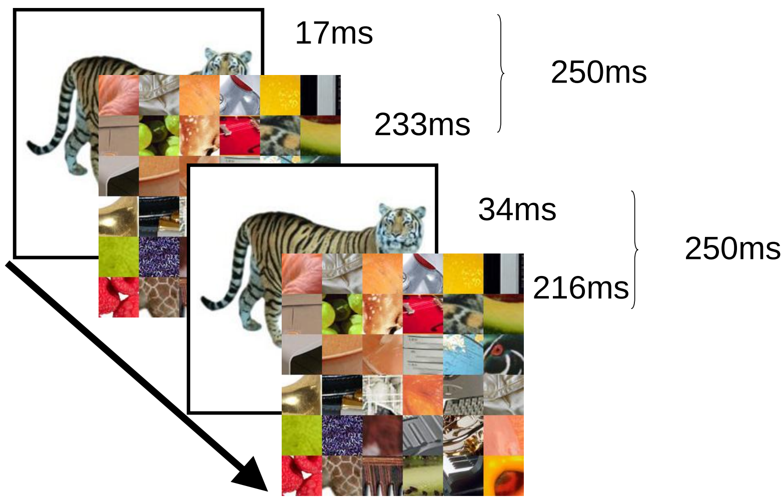

Figure 6.2: Timings of two blocks of the continuous identification procedure. The stimulus (tiger) is presented for 17 msec longer, and the mask for 17 msec less, in each successive block.

Results

Test Phase: Recognition

Recognition responses given on the 6-point scale were collapsed into *old* responses (items given a rating of 4-6) and *new* responses (ratings of 1-3) and are presented in Figure 6.3. A within-subjects ANOVA showed a significant effect of item type (seen/word/new) on recognition, $F(1.44, 27.44) = 240.08$, $MSE = .02$, *partial* $\eta^2 = .93$, $p < .001$ (Mauchly's test indicated that the assumption of sphericity had been violated for this main effect: $\chi^2(2) = 8.75$, $p < .05$; the Greenhouse-Geisser correction ($\varepsilon = .72$) was used). The proportion of hits to pictures which had appeared in the study phase ($M = .86$) as well as false alarms to pictures which had appeared as words in the study phase ($M = .21$) were both significantly higher than the proportion of false alarms to new pictures ($M = .14$; $t(19) = 17.52$, $SEM = .04$, $r = .94$, $p < .001$ and $t(19) = 3.09$, $SEM = .02$, $r = .24$, $p < .01$, respectively). Unlike in all the previous experiments reported so far (except Experiment 4, in which criterion was zero by definition), participants responded to seen and new items without any bias: the criterion in the seen-new comparison ($M = –.01$) was not significantly different to zero ($p > .9$; see Appendix B).

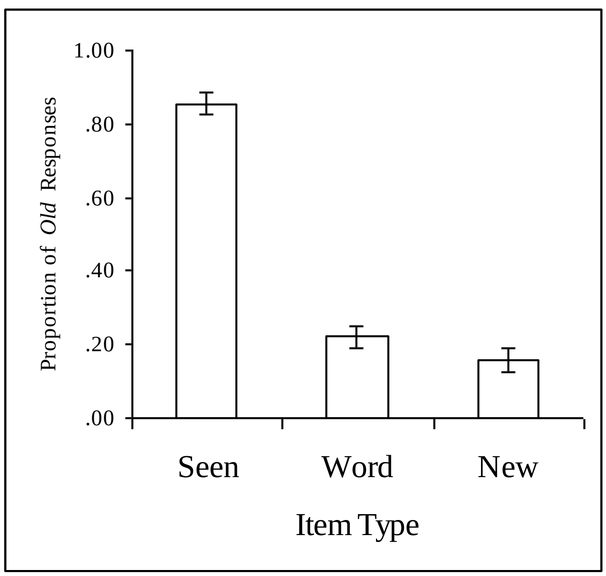

Figure 6.3: Mean proportion of old *responses by item type on the direct test in Experiment 5.*

The same pattern of recognition data emerged when only the high confidence (confidence level 6, *very sure old*) responses were analysed separately: a repeated-measures ANOVA showed a significant effect of item type, $F(1.15, 21.85) = 145.07$, $MSE = .03$, *partial* $\eta^2 = .88$, $p < .001$ (Mauchly's test indicated that the assumption of sphericity had been violated: $\chi^2(2) = 24.19$, $p < .001$; the Greenhouse-Geisser correction ($\varepsilon = .58$) was used). High confidence hits to seen items were made at a rate of .68, while high confidence false alarms to word items were made at a rate of .07, which was significantly higher than the false alarm rate of .04 to new items, $t(19) = 2.15$, $SEM = .02$, $r = .19$, $p < .05$; only seven participants made any high confidence false alarms to new items, while 13 of the 20 participants made high confidence false alarms to word items. The average confidence of *old* responses was compared between the three conditions for those 10 participants (half of the sample) who made at least 4 *old* responses in each category. A within-subjects ANOVA confirmed that confidence levels differed between conditions, $F(2, 18) = 25.40$, $MSE = .09$, *partial* $\eta^2 = .74$, $p < .001$. While hits to seen items were made with higher confidence (M = 5.69) than false alarms to new items ($M = 4.78$; $t(9) = 5.68$, $SEM = .16$, $r = 71$, $p < .001$), confidence levels did not differ significantly between new item and word item false alarms ($M = 4.98$; $t(9) = 1.86$, $SEM = 0.11$, $r = .18$, $p = .10$).

Test Phase: Priming

Individuals' median RTs by item type were used in the analyses. RTs were excluded where the response given did not match the object name selected by the experimenter to serve as the label for the associated picture, or any of its synonyms. Identification accuracy was high (.94) across all conditions; a within-subjects ANOVA indicated that accuracy differed between conditions, $F(2, 38) = 4.55$, $MSE = .01$, *partial* $\eta^2 = .19$, $p < .05$.

Participants were more accurate in identifying seen items (.96) than new items (.94; $t(19) = 2.40$, $SEM = .01$, $r = .29$, $p < .05$), and equally accurate in the word (.93) and new conditions ($p = .25$). Of the 1800 trials across all participants, 103 trials were excluded due to inaccurate identification. RTs were also excluded on the 7 trials where the RT exceeded the length of the CID cycle (3500 msec). Overall, these exclusion criteria resulted in a loss of 6.1% of the data. Figure 6.4 shows median RTs in milliseconds by item type. Priming is defined by a decrease in RT compared to baseline performance on new items.

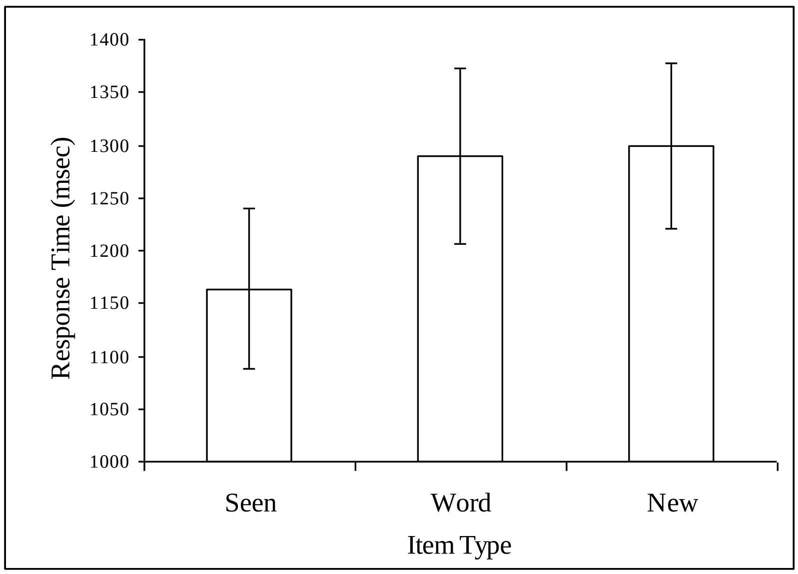

Figure 6.4: Median perceptual identification RTs by item type on the indirect test in Experiment 5.

Item type (seen/word/new) had a significant effect on RTs, $F(2, 38) = 11.38$, $MSE = 10{,}087$, $partial\ \eta^2 = .38$, $p < .001$. However, while priming of seen items was significant, $t(19) = 4.13$, $SEM = 33$, $r = .19$, $p < .01$, there was no significant priming of pictures which had appeared as words in the study phase ($t < 1$)[3]. Furthermore, only 9/20 participants showed the effect in the required direction (i.e., they were on average faster in responding to studied word items than new items, allowing for 17 msec measurement error, the equivalent of one screen refresh). A paired-samples *t* test confirmed that priming differed significantly between the studied picture and studied word conditions, $t(19) = 3.84$, $SEM = 33$, $r = .41$, $p < .001$ (see Appendix C for a summary of recognition and priming results for Experiments 5–7).

Discussion

In this experiment, pictures that had been studied were correctly recognised at a high rate in the direct test and were also identified faster than new pictures on the indirect test. Participants also made more false alarms to pictures that had appeared in the study phase as words than to new pictures. However, pictures that had appeared in the study phase as words were identified no faster on the indirect test than were new pictures. This provides evidence that the continuous identification task when applied to pictures is sensitive to variations in visual features between study and test, shows no detectable cross-form priming, and hence is perceptually driven (Jacoby & Hayman, 1987).

In Experiments 6 and7, false memories created by the imagination manipulation used in previous experiments were tested with this indirect task. Imagining a picture may have an equivalent (if lower) effect on perceptual processing to actually seeing one. McDermott and Roediger (1994) found that imagery boosted word to picture priming on a test of perceptual identification of recoverable picture fragments, although once again a pure list design was used. Therefore, priming of imagined items might occur in our experiment if the process of imagining an item has an effect similar to that of perception. Alternatively, priming for these pictures may arise from prior presentation of their word label due to nonperceptual processes as found by Park and Gabrieli (1995).

Whether participants show priming after imagination will depend to some extent on the perceptual overlap between the images they create and the test stimuli. Biederman's (1987) recognition-by-components theory proposes that priming relies on the overlap of geometric components (geons) between study and test. According to this theory, perceptual priming will be preserved despite differences between study and test stimuli, as long as these changes do not affect the structure of the objects (Biederman & Cooper, 1991a). Manipulations of object size (Biederman & Cooper, 1992) and viewing angle (Biederman & Gerhardstein, 1993) have confirmed that this is the case. More relevant to our research, Bartram (1974) and Biederman and Cooper (1991b) showed that naming an object facilitates naming another object (or token) corresponding to the same label (or type), though to a lesser extent than naming the same object. A representation created through imagery can be seen as another token of the type embodied by the associated object label. The few studies that have looked specifically at whether imagery can lead to perceptual priming have produced mixed results. McDermott and Roediger (1994) found priming in a perceptual task which involved the identification of pictures from fragments presented for 100 msec and 200 msec following imagination of these pictures from their word labels. Michelon and Koenig (2002, Experiment 2), on the other hand, only found significant priming on an object identification task after perception and not after imagery.

These experiments investigate whether false recognition of pictures could in part be attributed to perceptual overlap between the test stimulus and a representation that has been formed as a result of an experience other than perception. If this were case, we would expect participants to show both false recognition and perceptual priming of imagined items.

6.2 Experiment 6

In this experiment, false memory was induced using the same imagination inflation procedure as in Experiments 1a, 1b, 2, and 4, and tested with the combined continuous identification and recognition procedure employed in Experiment 5.

Method

Participants

Fifty-one Undergraduate students from University College London (6 males) 18-29 years of age (*M* = 18.8) participated in the experiment as part of a course requirement.

Design

With respect to manipulated variables, the experiment was identical to Experiment 5: A repeated measures design was used with item type (seen/imagined/new) as a within-subjects variable, and two dependent measures, proportion judged as *old* in the recognition test and RT to identify pictures in the priming test.

Procedure

The study and imagery phases were similar to those of previous experiments: in the study phase, participants labelled pictures of objects, while in the imagery phase, they studied the objects' names and tried to recall whether the associated objects had appeared in black and white or colour in the study phase. The imagery task was the one used in experiments 1a, 1b, 2, and 4. As in previous experiments, ten of the object names were replaced with the ten critical imagined items. The number of items in each phase corresponded to that in Experiment 1-3. The number of times participants encountered the critical imagined items was consistent with Experiments 3 and 4, in which participants encountered the imagined items until the imagery phase, where they encountered them once only. This was done in order to be sure that the false recognition and any potential priming effect were a result of this one interaction with the imagined items. The test phase of the present experiment was identical to that of Experiment 5: Participants identified a briefly flashing picture on each trial and then indicated whether this pictures was *old* (i.e., whether they had seen the picture in the study phase) or *new* on a confidence scale from 1 to 6 where 1 = *very sure new* and 6 = *very sure old*. As in previous experiments, participants were instructed specifically to use the old response only when they had seen the picture itself earlier in the experiment.

Results

Study and Imagery Phases

Performance in the study phase ranged from 88.9% to 100% correct identification of pictures from a choice of two labels, with a mean score of 97.5%. In other words, over half the participants made less than 3 identification errors in the 90 trials, and no-one made more than 10 errors. Performance in the imagery phase was based on the 80 items which had appeared in the study phase; the 10 imagined items were excluded from the analysis as there was no correct response possible. In this task, participants had to indicate whether they had seen each picture in black and white or in colour in the study phase. Scores ranged from 52.5% to 83.8%, with a mean score of 65.7% (SD = 8.2), which was significantly different from the chance level of 50%, $t(50) = 13.61$, $p < .001$. On the 1 to 5 confidence scale, participants were significantly more confident on their answers to the 80 items they had actually seen as pictures (M = 3.41) than to the 10 imagined items for which there was no correct response because these items had not in fact been seen, (M = 2.00; $t(50) = 12.81$, *SEM* = .11, *r* = .67, *p* < .001). As in previous experiments, performance on this task was consistently above chance, indicating that participants were paying attention to the task and following instructions.

Test Phase: Recognition

For simplicity, responses given on the 6-point scale were collapsed into *old* responses (items given a rating of 4–6) and *new* responses (items rated 1–3). Figure 6.5 shows the proportion of *old* responses to seen, imagined, and new items.

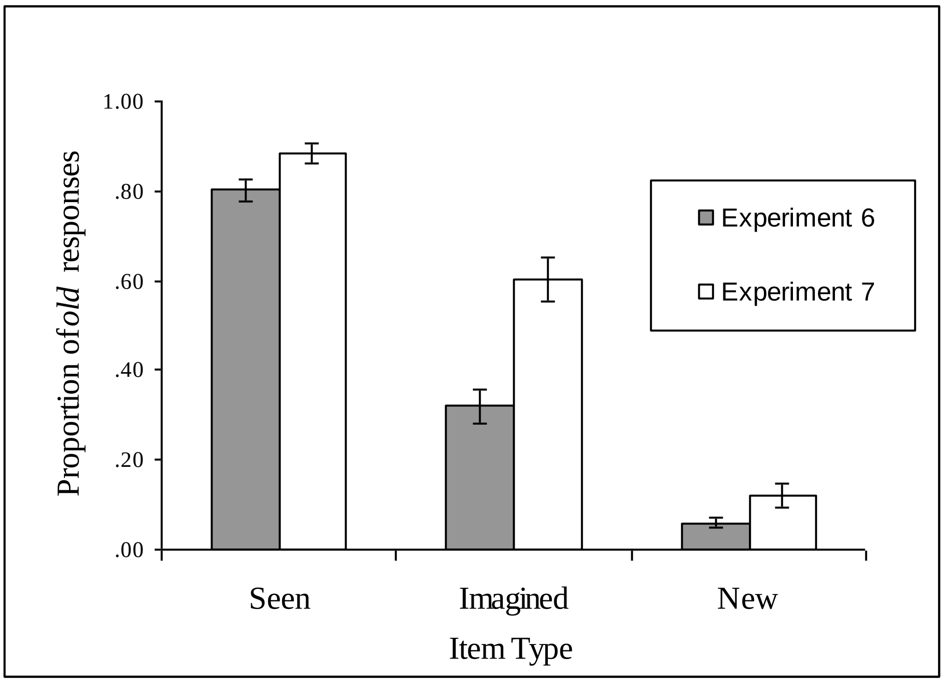

Figure 6.5: Mean proportion of old *responses by item type on the direct test in Experiments 6 and 7.*

A repeated measures ANOVA showed that there was a significant effect of item type, $F(1.61, 80.2) = 233.35$, $MSE = .04$, *partial* $\eta^2 = .82$, $p < .001$ (Mauchly's test indicated that the assumption of sphericity had been violated for the main effect of the within-subjects variable, item type: $\chi^2(2) = 13.9$, $p < .01$; a Greenhouse-Geisser estimate of sphericity was used to correct degrees of freedom, $\varepsilon = .80$). The likelihood of participants responding *old* to an item was therefore dependent upon its status. The proportion of hits to pictures which had appeared in the study phase ($M = .80$) as well as false alarms to pictures which had appeared as words in the imagery phase ($M = .32$) were both significantly higher than the proportion of false alarms to new pictures ($M = .06$; $t(50) = 28.41$, $SEM = .03$, $r = .94$, $p < .001$ and $t(50) = 7.41$, $SEM = .04$, $r = .55$, $p < .001$ respectively), and 40 of 51 participants showed the latter effect (i.e., their imagined item false alarm rate was at least one item higher than their new item false alarm rate).

For high confidence responses, a repeated-measures ANOVA showed a significant effect of item type, $F(1.47, 73.47) = 207.87$, $MSE = .03$, *partial* $\eta^2 = .81$, $p < .001$ (Mauchly's test indicated that the assumption of sphericity had been violated: $\chi^2(2) = 21.95$, $p < .001$; the Greenhouse-Geisser correction ($\varepsilon = .74$) was used). High confidence hits to seen items were made at a rate of .57, while high confidence false alarms to studied word items were made at a mean rate of .09, significantly higher than the false alarm rate of .03 to new items, $t(50) = 3.68$; $SEM = .02$, $r = .31$, $p = .001$; only 11 participants made any high confidence false alarms to new items, while 26 of the 51 participants made high confidence false alarms to studied word items. Due to the low number of false alarms to new items, average confidence of *old* responses was compared between the seen and imagined conditions only for those 23 participants who made at least 4 *old* responses in each of the two categories. A paired-samples t-test confirmed that hits to seen items were made with higher confidence than were false alarms to studied word items ($M = 5.60$ and $M = 4.90$ respectively; $t(22) = 9.28$, $SEM = .08$, $r = .61$, $p < .001$). As in Experiments 1a and 1b, the criterion in the seen-new comparison, (*M* = .25), was significantly stricter than zero, $t(50) = 5.45$, $p < .001$[1] (see Appendix B).

Test Phase: Priming

The same exclusion criteria were applied to RTs as those used in Experiment 5. Identification accuracy was high (.95) and did not differ between conditions ($p > .4$). Of the 1530 trials across all participants, 70 trials were excluded due to inaccurate identification, and 5 trials were excluded because the RT exceeded the length of the CID cycle (3500 msec). Overall, these exclusion criteria resulted in a loss of 4.9% of the data. Figure 6.6 shows median RTs in milliseconds by item type.

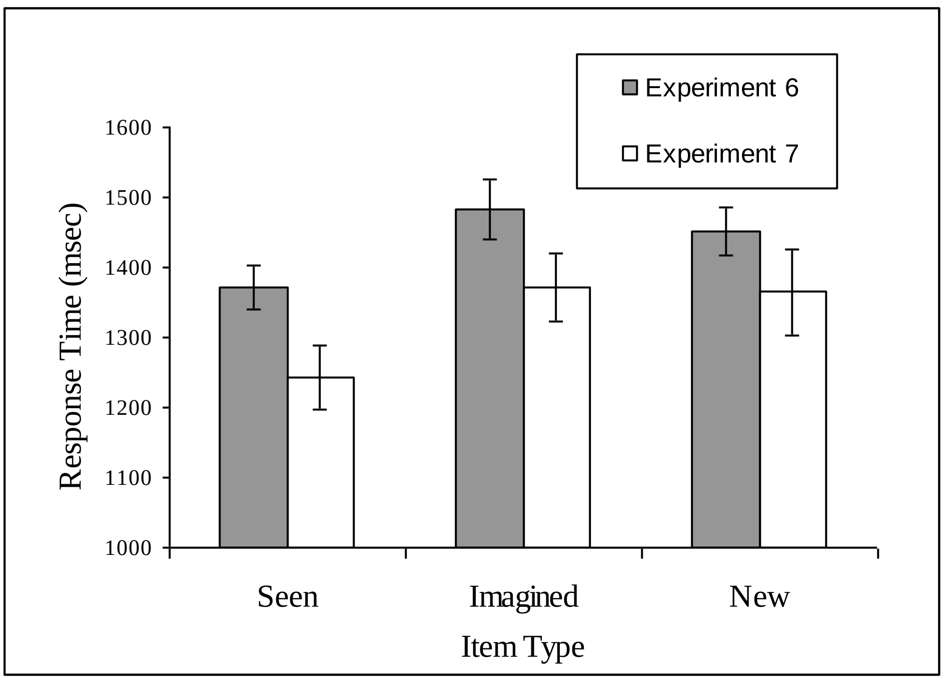

Figure 6.6: Median perceptual identification RTs by item type on the indirect test in Experiments 6 and 7.

A repeated measures ANOVA performed on seen/imagined/new RTs produced a significant effect of the within-subjects variable, item type, on perceptual identification,

$F(2, 100) = 9.12$, $MSE = 17{,}977$, *partial* $\eta^2 = .15$, $p < .001$. Thus, RTs were dependent upon the item’s status. Planned comparisons revealed that while seen items were recognised significantly faster than new items, $t(50) = 3.03$, $SEM = 25$, $r = .16$, $p < .01$, there was no significant difference in RTs between imagined and new items, $t(50) = –1.18$, $SEM = 29$, $r = –.07$, $p = .21$[4]. In fact, the numerical difference was in the opposite direction to that predicted, and only 20 of the 51 participants showed the pattern of results that would be indicative of priming (as in Experiment 5, the criterion was a minimum difference in RTs of 17 msec). A paired-samples *t* test confirmed that priming differed significantly between the seen and imagined conditions, $t(50) = 4.26$, $SEM = 26$, $r = .28$, $p < .001$ (see Appendix C).

Discussion

This experiment, unlike Experiment 5, included a false memory induction with an imagery component. As this component increases false alarms to pictures that are encountered in word form, we might have predicted that it would also lead to priming of these items on the perceptual identification test. However, no such priming was found. The absence of priming for pictures of imagined objects on a perceptual identification task is in line with the Michelon and Koenig (2002) finding of no perceptual priming following imagination, the Hicks and Starns (2005) finding of no priming for critical lures in the DRM paradigm on a verbal perceptual identification task, and the studies that found no priming for critical lures on a lexical decision task (McKone, 2004; Zeelenberg & Pecher, 2002). The findings are also in line with studies which have found no cross-form priming from words to pictures (Hirshman et al., 1990; Warren & Morton, 1982; Weldon et al., 1995). However, results of Experiment 6 contradict McDermott and Roediger (1994) who found priming in a perceptual task which involved the identification of pictures from fragments presented for 100/200ms following imagination of these pictures from their word labels.

6.3 Experiment 7

It is possible that the priming effect is weaker than the false recognition because priming responses are intrinsically noisier than direct measures (Buchner & Brandt, 2003; Berry, Henson, & Shanks, 2006). In order to test this possibility, it was necessary to use a false memory manipulation that produced an even higher level of false recognition. Goff and Roediger (1998) examined false recognition of imagined actions and found that increasing the number of times participants imagined performing an action lead to an increased tendency to report having performed the action. Following this rationale, higher levels of false recognition are expected in the procedure employed in Experiments 1 and 2, where participants were exposed to imagined items in the study phase as well as the imagery phase compared with that in Experiments 3 and 4 where they encountered them once in the imagery phase. This stronger false memory induction procedure was used in the present experiment in an attempt to create optimal conditions for detecting a priming effect.

Method

Participants

Twenty-four volunteers (8 males) 18–43 years of age (M = 23.0 years) were recruited from within University College London to participate in the experiment and were reimbursed for their time.

Procedure

The procedure was almost identical to that of Experiment 6, except for the number of times participants encountered the critical imagined items. This was increased from once in Experiment 6 (in the imagery phase only) to four times in the current experiment (once in the imagery phase and three times in the study phase) as incorrect options on the labelling task to match Experiments 1 and 2. The imagery phase was identical to that of Experiment 6. Finally, the test phase was also identical to that of Experiment 6, except that on the direct test participants had to simply indicate whether a picture was *old* or *new* instead of making a 1–6 confidence rating.

Results

Study and Imagery Phases

In the study phase, participants identified images from a choice of two words with a mean correct score of 97.0% (*SD* = 4.65). Performance in the imagery phase, where participants reported the colour format of pictures seen in the study phase, ranged from 55.0% to 83.4%, with a mean score of 66.7% (*SD* = 8.9), which was significantly above the chance level of 50%, $t(23) = 9.16$, $p < .001$. As in Experiment 6, in the imagery phase participants were significantly more confident on the 5-point scale in their answers to the 80 pictures from the study phase ($M = 3.56$) than to the 10 imagined items ($M = 2.59$; $t(23) = 7.01$, $SEM = .14$, $r = .59$, $p < .001$). In a between-experiment comparison of Experiments 6 and 7, while participants were equally confident in their answers to items from the study phase ($p = .39$), participants in the current experiment were more confident in their answers to the imagined items, $t(73) = 2.99$, $SEM = .20$, $r = .35$, $p < .01$. This difference can be attributed to the fact that participants in the current experiment had previously encountered these novel items as words in the study phase. Aside from this, there were no significant differences in performance in any of the above measures between Experiment 7 and Experiment 6.

Test Phase: Recognition

Figure 6.5 shows the proportion of *old* responses to the three item types. A repeated measures ANOVA showed that there was a significant effect of item type, $F(1.58, 36.3) = 166.01$, $MSE = .03$, *partial* $\eta^2 = .88$, $p < .001$ (Mauchly's test indicated that the assumption of sphericity had been violated for the main effect of the within-subjects variable, item type: $\chi^2(2) = 6.85$, $p < .05$; a Greenhouse-Geisser estimate of sphericity was used to correct degrees of freedom, $\varepsilon = .79$). The proportion of hits to pictures which had appeared in the study phase ($M = .88$) as well as false alarms to pictures which had appeared as words in the study and imagery phases ($M = .60$) were both significantly higher than the proportion of false alarms to new pictures ($M = .12$; $t(23) = 22.82$, $SEM = .03$, $r = .95$, $p < .001$ and $t(23) = 12.15$, $SEM = .04$, $r = .78$, $p < .001$ respectively), and all 24 participants showed the latter effect. Unlike in Experiment 6, participants responded to seen and new items without any bias: the criterion in the seen-new comparison, ($M = –.00$), was not significantly different to zero ($p > .9$; see Appendix B). Participants did not differ in bias between Experiments 5 and the current experiment ($p > .7$).

Test Phase: Priming

Individuals' median RTs by item type were analysed with the same exclusion criteria as in Experiments 5 and 6. Identification accuracy was very high (.97) and did not differ between conditions; of the 720 trials across all participants, 30 trials (4.2%) were excluded due to inaccurate identification. Figure 6.6 shows median RTs in milliseconds by item type. A repeated measures ANOVA performed on seen/imagined/new item RTs produced a significant effect of the within-subjects variable, item type, $F(2, 46) = 7.16$, $MSE = 17{,}660$, *partial* $\eta^2 = .24$, $p < .01$. Planned comparisons revealed that while seen items were recognised significantly faster than new items, $t(23) = 3.08$, $SEM = 40$, $r = .23$, $p < .01$, with 21 of 24 participants showing the effect, there was no significant difference in RTs between imagined and new items ($t < 1$)[5]; once again, the numerical effect was in the opposite direction to that expected if performance on the indirect test mirrored that of the direct test, and only 11 of 24 participants showed the effect. A paired-samples *t* test confirmed that priming differed significantly between the seen and imagined conditions, $t(23) = 4.03$, $SEM = 32$, $r = .30$, $p < .001$ (see Appendix C).

Discussion

The current experiment yielded both the highest rate of false recognition of imagined items (.60, almost double the next highest rate of .37 achieved in Experiment 1b) and the highest corrected false recognition rate (.48, compared with .31 in Experiment 1b) of all experiments reported in the thesis. Three factors may have contributed to the effect. Firstly, as in Experiments 1 and 2, but unlike in Experiments 3, 4, and 6, participants encountered the imagined items four times. As in Goff and Roediger's (1998) imagination inflation study, this increased the tendency to confuse imagination with perception. Secondly, the current experiment was the only one in which the recognition test involved a simple old/new judgment. According to the dual-process theory (e.g., Yonelinas, 1994), familiarity and recollection processes differentially affect memory performance. A more straightforward task such as old/new recognition (when compared with having to make judgments on a confidence scale) may have allowed participants to make quicker responses and thus rely more on familiarity than recollection. This would be expected to both relax the response criterion because all items came from familiar categories, and increase false alarms to imagined pictures specifically. Finally, Experiments 6 and 7 differed from other experiments in that each recognition trial was preceded by perceptual identification. The revelation effect (Watkins & Peynircioglu, 1990) describes the finding that viewing a slowly revealing item at test leads to an increased tendency to call that item ‘old’. It has generally been found that the effect is greater for new relative to old items (e.g. Westerman & Greene, 1996), suggesting that the revelation manipulation affects memory sensitivity and leads to a decrease in accuracy. Although it is not clear what effect such a manipulation would have on previously imagined items, the perceptual identification component is likely to have led to participants adopting a more relaxed response criterion in Experiments 6 and 7. Indeed, participants in Experiment 7 performed the recognition task with the least strict criterion when compared with all other experiments using a similar procedure (except for Experiments 4 and 5).

While the changes made from Experiment 6 and Experiment 7 almost doubled both the absolute and corrected false recognition rates, there was no analogous effect on priming of imagined items. As in Experiment 6, these items were recognised no faster than new items.

6.4 Discussion of Experiments 5–7

In Experiments 5-7, we showed a consistent false memory effect for studied word items in the direct test and a consistent lack of priming of these items on the perceptual identification test. Figure 6.3 shows that while hits to studied picture items and false alarms to new items remained stable between the three experiments, the false alarm rate to studied word items depended on the procedure and, consistent with Goff and Roediger's (1998) imagination inflation study, was highest when participants were asked to imagine items four times. This false alarm rate was always significantly above that of new items. Figure 6.4, on the other hand, shows that the procedure did not affect how fast studied word items were identified relative to new items; seen items were always identified significantly faster than new items, while imagined items were identified at roughly the same speed as new items. Note that participants performed the identification task at different speeds in the three experiments. Experiment 5 yielded the fastest overall RTs, and this is most probably a result of participants speeding up across the 90 items in the test phase; in Experiments 6 and 7 the test phase consisted of only 40 items consistent with all previous experiments in the series with the exception of Experiment 4. Participants also performed the task faster in Experiment 7 than in Experiment 6, and this could be because the recognition decision made on each trial (*old*/*new* as opposed to a 1–6 confidence judgment) required less time and effort and thus interfered less with the identification task. The important result, however, is that despite these variations in speed, the pattern of RTs was consistent across all three experiments. Appendix C presents a comparison of performance on the direct and indirect tests in Experiments 5–7. It is clear from these data that while false recognition of studied word items increased with imagery instructions, no manipulation produced any perceptual priming of these items.

The lack of priming for studied word items on the perceptual identification task in Experiments 6 and 7 suggests that there is little perceptual overlap between the representations of falsely recognised pictures and the actual stimuli in the current procedure. Nonetheless, false recognition of these items is high following one imagination event (Experiment 6; see also Experiments 3 and 4) and even higher following multiple imagination events (Experiment 7; see also Experiments 1a, 1b, 2, and 4). Of course, we cannot be certain that participants did in fact follow instructions to imagine objects. However, the imagery task fulfilled its function of inflating the false alarm rate on the direct test by increasing participants' interaction with visual representations of the objects; whatever the specific mental processes participants engage in during imagery, the conclusion that false recognition on the direct task is largely unaccompanied by perceptual priming remains unaffected. On the other hand, the pattern of results does not preclude the possibility that participants do create perceptually rich representations during imagery, but does suggest that something other than perceptual overlap between internal representations and test stimuli must be driving the false memory effect. In fact, any mismatch between the two ought to make it easier for participants to reject unstudied items. If this were the case, then we would expect a lower rate of false recognition when objects were imagined multiple times, strengthening the internal representations and making more salient any differences between them and the presented pictures. Instead, we see an increase in the false alarm rate, suggesting that nonperceptual processes are responsible for the effect.

McKone (2004) has suggested that perceptual indirect tests can help distinguish between true and false memories; more specifically, priming on indirect memory tests that are thought to be driven largely by perceptual processing should only occur for items that are actually studied. Experiments 5–7 add further support to such a view: priming on a perceptual identification test was only found for pictures that were actually studied, but not for pictures that had appeared in word form (Experiment 5), nor for pictures that were imagined once (Experiment 6) or imagined multiple times (Experiment 7). This absence of priming was contrasted with the presence of an increasingly large false recognition effect. This suggests that the representations of the falsely recognised items are dissimilar in nature to those of truly perceived items. Such an interpretation stems from the view that the direct and indirect tests tap a common memory representation. Of course, another possibility is that they are based on distinct memory systems (Squire, 1994).

CHAPTER 7

GENERAL DISCUSSION

This thesis makes three contributions, two of which are methodological. First, a new false memory induction procedure was developed which elicits high levels of false recognition of pictorial stimuli on three different direct memory tests: recognition (Experiments 1a and 1b), source monitoring (Experiments 2 and 3), and two-alternative forced choice (Experiment 4). Second, false memory for pictorial stimuli was examined on an indirect test (Experiments 5-7), a technique previously only employed with word stimuli. Finally, these methodological advances were applied to a theoretical question: specifying whether false memories which arise from imagination are predominantly perceptual or conceptual in nature.

7.1 False Memory for Pictures on Direct Tests

Experiments 1a and 1b demonstrated the basic procedure employed throughout the thesis, which involved three phases: a study phase where participants were presented with pictures in a labelling task, an imagery phase where they studied object names and were asked to think back to the pictures representing these objects, and a test phase where they were presented with pictures of studied objects, pictures of objects whose names had been studied, and new pictures from the same semantic categories. In these first two experiments, the test involved indicating whether each picture had appeared in an earlier part of the experiment. In both experiments, participants consistently made more false alarms to pictures whose names had been studied compared with new pictures. This was a robust effect which occurred in 51 of the 60 participants who took part in those initial studies. Although these experiments alone are striking in that they show higher levels of corrected false recognition of pictures than any other immediate test conditions in the literature, further experiments confirmed that this effect persisted under a variety of testing conditions.

One issue of note is that the false memory induction involved a mild form of deception: the imagery instructions stated that all items in the imagery phase had appeared as pictures in the study phase. It is possible that participants could have been aware of the false memory manipulation. In other words, they may have noticed that the some items in the imagery phase had not actually appeared as pictures in the study phase. Awareness of the deception could lead to two opposing influences on the final test: either, participants could be less likely to falsely recognize a new item because they would be able to reject any items they had encountered in the imagery phase that they could not remember studying in pictorial form; or, participants may be more likely to make false alarms to imagined items in order to conform with experimenter expectations. Unfortunately, the procedure did not involve a measure to assess this awareness; i.e., participants were not systematically asked whether they had noticed any items in the imagery phase that had not also appeared in the study phase. However, informally, only a very few participants mentioned that they were unsure of whether they had actually seen pictures of all of the items in the imagery phase, and those who did were very unconfident in this assertion. Furthermore, in Experiments 2 and 3 participants were explicitly told that some items had only appeared in word form, so that participants could demonstrate any awareness of the manipulation.

In Experiments 1a and 1b, participants were instructed only to respond ‘old’ to pictures that they had actually seen in the study phase. In Experiment 2, on the other hand, participants were additionally instructed to classify pictures that they had only imagined into a separate ‘imagined only’ category. It could be argued that Experiments 1a and 1b masked any knowledge participants might have had about the true source of imagined pictures, and that given a response option that directly matched this source, they would be able to use it accurately. However, even with this additional response option, participants still indicated that they had actually seen more of the imagined pictures than the new pictures. This finding concurs with previous studies showing that imagination leads to an increased likelihood to claim that a picture had been studied (e.g., Henkel & Franklin, 1998), and, crucially, extends the effect to a test in which the pictures themselves are presented as memory probes. Experiment 3 replicated this effect with a manipulation that shed light on the theoretical question regarding the locus of these false memories, which is discussed below (see Section 7.3, p. 106).

Whilst the false memory effect remained in the source experiments, it was significantly reduced. An attributional account can explain this reduction. Hoffman (1997) showed that when people make false alarms, these are most often attributed to the memory source that is judged to be the weakest. While perception is deemed to be a weaker source than imagination for words (Hoffman, 1997, Experiment 1; “I don’t remember imagining it, so it had to be you”), for pictures imagination is deemed to be a weaker source than perception (Hoffman, 1997, Experiment 2: “I don’t remember seeing it, so it had to be me”). In the current procedure, if a previously imagined picture feels slightly familiar, participants may correctly classify it as *imagined* instead of *perceived and imagined* not because they have accurate source memory, but because that is the weaker of the two potential memory sources. Thus, using strength as a basis for source judgments, pictures that evoke stronger memories would be judged as *perceived and imagined*, whereas pictures that invoke weaker memories would be judged as *imagined*.

When hit and false alarm data from recognition tests are presented in the context of false memory research, it is often implied if not overtly stated that in a situation where the false alarm rate is high, participants might be unsure which of two items had been studied if they were presented with one old and one new item. Alternatively, an ‘old’ response to a studied item may be made on different grounds to an ‘old’ response to a new item, and participants may be aware of using such a strategy. If this is the case, in a situation where participants were forced to choose between a studied item and a new item, they should have no difficulty indicating which item had been studied. Of course, collecting additional data such as remember/know and confidence ratings (as was done in Experiments 1a and 1b respectively) can shed some light on this issue, but a more definitive method would be to pit studied items directly against new items in a forced-choice test. If participants are aware of the difference between studied and new items which they call ‘old’, they should perform this task at ceiling and, more crucially, it should not matter whether a studied item is paired with a distractor that typically elicits high rates of false alarms, or a distractor that is easily rejected. This prediction is difficult to test in the DRM paradigm because due to the nature of the list design, studied and critical items cannot be counterbalanced. This leads to two potential confounds: firstly, studied and critical items differ inherently on a variety of measures other than whether they are presented to participants at study; and secondly, critical items are inherently more similar to studied items than are unrelated items.

Although these issues can be overcome (see McDermott et al., 2009), the procedure presented here affords a more straightforward design. In Experiment 4, participants were tested on picture pairs in which studied pictures were paired either with a new picture or with a picture that had been imagined. Crucially, all items came from the same semantic categories and were entirely counterbalanced between conditions, so any inherent inter-item similarities and differences were controlled for. Previous studies have shown that people are very good at this sort of task (Shepard, 1967), and indeed, in the experiment reported here participants were almost perfectly accurate in selecting the studied picture when the alternative was a new picture. However, in concordance with the higher false alarm rate for imagined compared to new pictures in Experiments 1a and 1b, participants in Experiment 4 found it harder to correctly identify the studied picture from a pair when the distractor had been imagined.

In the analysis of Experiment 4, the focus was on how likely participants are to recognise the seen item as having been studied in the seen-new and seen-imagined pairs, and thus correctly select it from the two options. An alternative story could be told about the item in each pair that had not been studied. In other words, it is possible that instead of looking for the familiar item (or the item with the strongest signal) to make their selection, at least in some cases participants may look for an item that jumps out as unfamiliar and then select the alternate item in the pair. This explanation does not require there to be any memory for seen items: a correct response can be made purely on the basis of being sure that the item it is paired with was not studied. The fact that participants chose the imagined item more often than the new item on imagined-new trials shows that the imagined items were indeed more familiar than new items. Consequently, if participants were to use a rejection strategy, this would be less successful for seen-imagined than seen-new trials. However, such a strategy alone cannot account for the results, as it would predict equivalent likelihood of rejection of the new item in seen-new and imagined-new pairs, whereas in fact, participants were able to reject the new pair in 96% of seen-new trials and only in 72% of imagined-new trials. Hence, the difference in accuracy on seen-imagined compared with seen-new pairs is better described by a signal detection account. According to this account, decisions are based on the comparison of the signal strength of the two items in a pair, and the decreased accuracy on seen-imagined compared with seen-new pairs arises from the fact that imagined items exceed seen item in terms of strength more often than do new item.

The procedure proposed in this thesis (and Weinstein & Shanks, 2008) fills the void that existed in the field with respect to false memory procedures involving recognition of detailed, distinctive pictures. While other procedures (e.g., the DRM paradigm, Roediger & McDermott, 1995; the category associates procedure, Koutstaal & Schacter, 1997) have dealt with false recognition arising from studying semantically related associates, the procedure presented here deals specifically with the effect of interacting with a concept (i.e., encountering the name of an object) on recognition memory of pictures. In other words, it was designed to demonstrate that encountering an item conceptually (i.e., as a word) could lead to this item being incorrectly identified as having been studied even though participants were basing their recognition decision on a highly perceptually distinctive form (i.e., as a picture). This finding can be readily applied to situations outside the laboratory, including legal contexts, where it would have strong implications in terms of the reliability of recognition memory for physical evidence that had previously been alluded to verbally.

7.2 False Memory for Pictures on an Indirect Test

Indirect tests have been used to study false memory for verbal stimuli (e.g., McDermott, 1997; see Chapter 5), while others have investigated whether priming of pictorial stimuli on perceptual indirect tests can arise from imagery alone (e.g., Michelon & Koenig, 2002), a hypothesis which is based on the premise that imagery and perception share processes (Kosslyn, Thompson, & Alpert, 1997) and potentially even produce comparable representations (Farah, Péronnet, Gonon, & Girard, 1988). Experiments 5-7 bridge the gap between these two fields of investigation that have until now remained separate, by applying the indirect test methodology to a false memory procedure involving pictorial stimuli.

In three experiments, there was no priming on a picture perceptual identification test following study of the names of the objects depicted in these pictures (Experiment 5), nor following visual imagery of these objects performed once (Experiment 6) or four times (Experiment 7). Conversely, false alarms to these items on the direct recognition test increased with visual imagery instructions, producing a dissociation between the two measures. With respect to the extant literature applying indirect perceptual tests to verbal stimuli, the above finding is consistent with Hicks and Starns (2005) who found no priming of critical items in the DRM paradigm on a similar task involving identification of words. Hicks and Starns claimed that no priming was found on this task because it relied more heavily on perceptual processes than any other indirect memory task; the identification task used here is the analogous task for pictorial stimuli, so the same explanation applies. Turning to the imagery literature, some previous findings (Cabeza, Burton, Kelly, & Akamatsu, 1997; McDermott & Roediger, 1994) may have led to the prediction that perceptual priming should occur for imagined items in the current procedure. However, the lack of priming is consistent with Michelon and Koenig's (2002, Experiment 2) finding of no perceptual priming of objects following imagination.

Looking at Michelon and Koenig's (2002) results more closely may help to elucidate the discrepancy between studies that have found priming from imagination (e.g., McDermott & Roediger, 1994) and the lack of priming reported here. In a between-subjects experiment (Michelon & Koenig, Experiment 1a), participants were either presented with photographs of objects, or presented with object names and asked to imagine the associated picture. Both groups were asked to count the number of parts in their photograph or image. At test, both participants were presented with pictures of parts of objects they had studied or imagined and parts of new objects for 67 msec each, and asked to identify the object. With proportion correct identifications as the dependent measure, participants were better at identifying objects from a briefly presented part after studying or imagining whole objects compared with new objects. Crucially, when the test phase involved identifying objects from briefly presented photographs of the entire object instead of from object parts, priming was only found after perception and not after imagination of the object from its referent (Michelon & Koenig, Experiment 2). The authors explain this result as arising from the top-down processing demanded by the more difficult task of identifying objects from their parts as compared with identifying entire objects. The encoding and identification tasks used in Experiments 6-7 of are highly similar to those of Michelon and

Koenig’s Experiment 2, with the exception of list design (pure versus mixed): in both cases, participants perceive and/or imagine entire objects and then identify briefly presented pictures of entire objects. The results reported here are therefore consistent with Michelon and Koenig and could be explained by the low cognitive demands of the perceptual identification task. One might thus expect priming of imagined items in the current procedure to appear on a more cognitively demanding task.

7.3 Perceptual and Conceptual Processes in False Recognition of Pictures

Another way of explaining the above results is with reference to the perceptual/conceptual dichotomy. One hypothesis is that a memory of an imagined item, much like that of a perceived item, contains perceptual features that partially match those of the test stimulus. This view is supported by the findings of perceptual priming of pictures following imagination described above (e.g., McDermott & Roediger, 1994). The opposite hypothesis is that the relevant concept is activated, but no perceptual trace exists. Evidence for this view comes from studies which have reported no priming of critical lures on perceptual indirect tests (McKone, 2004; Hicks & Starns, 2005), studies which have found that participants recall more distinctive, item-specific perceptual details in relation to true memories than false memories (e.g., Norman & Schacter, 1997), and neuroimaging data indicating that false memories show similar activation to true memories in conceptual brain regions but dissimilar activation in perceptual regions (Cabeza, Rao, Wagner, Mayer, & Schacter, 2001.

In the experiments reported here, recognition can be seen as a conceptual task and perceptual identification as a data-driven or perceptual task. This distinction is made on the basis of the finding that changes in perceptual features affect priming but not recognition (Jacoby & Dallas, 1981). According to this logic, the absence of priming for imagined items would indicate that the underlying memory that leads participants to report having perceived a picture when they had actually only studied the associated word is conceptual in nature. In other words, there is a memory of having studied a concept, but there is no memory of having encoded a picture. A false memory of a non-presented item may also contain perceptual features which are different to those of the actual stimuli. This mismatch ought to serve as a cue to reject unstudied pictures; in the experiments reported here, however, the more times participants encountered object names, the more false alarms they made (Experiment 7 compared to Experiment 6), consistent with Goff and Roediger's (1998) imagination inflation procedure with action events, rendering this explanation unlikely. The absence of priming following imagination found in Experiments 6 and 7 accords nicely with the finding of Experiment 3: when participants were asked to imagine sentences in response to word labels (conceptual processing), at test they incorrectly claimed to have studied the associated picture as often as when they were asked to imagine pictures in response to the word labels (perceptual processing). But the major novel finding is that this conceptual activation is sufficient to cause people to misrecognize a distinctive, but unseen, picture.

Jacoby and Whitehouse's (1989) classic finding of increased false alarms following briefly presented primes points to the idea that people misinterpret perceptual fluency as an indication of prior perception. This idea generated the hypothesis that in the current procedure, false recognition of imagined items may arise from misattribution of perceptual fluency. Experiments 3 and 5-7, however, did not produce any evidence for this hypothesis, showing, respectively, that false recognition of imagined items was just as likely following conceptual imagery as following perceptual imagery, and that it was unaccompanied by priming on a perceptual identification task. However, it is possible that the false recognition effect reported here is driven by misattribution of fluency that is conceptual in nature.

7.4 Future Directions

This thesis proposes a new procedure for the rapid induction of high levels of false recognition for pictorial stimuli and explores one theoretical question: whether the false recognition effect is driven by perceptual processes. Other pertinent questions may also be answered by using this procedure: for instance, it could be used to delve further into the effects of distinctiveness on memory, an issue which has generated mixed results (see Chapter 1, Section 1.5, p. 14). While distinctiveness was not part of the design in the experiments reported here, this factor could be investigated by manipulating the typicality of items presented at study and at test, and observing whether the probability of falsely recognising an imagined item is reduced if the stimulus is an atypical representation of the object. Alternatively, some stimuli could be replaced by pictures of the same objects taken from a different angle. The effects of these manipulations on false recognition could also be compared to performance on the indirect perceptual identification task, which would be sensitive to some but not all these perceptual changes (see Discussion of Experiment 5, p. 86). Further work could also be carried out to determine the exact nature of the activation of the imagined items and how that leads to them being falsely recognised. For instance, although participants may not recognise the physical properties of a previously imagined picture any faster than those of a new picture, they may be faster at processing the meaning of the picture, just as they would be if they had studied the picture itself, and this fluency may be misattributed to prior perception.

For the last decade, research into false memory has clustered around the DRM paradigm. The popularity of this procedure stems from the fact that it is quick, easy to replicate, and produces robust levels of false memory. While it has proved to be a very useful methodology which has helped answer many questions about associative memory illusions, it is limited to answering questions about this particular type of false memory, and has seldom been applied to stimuli other than words (see Israel & Schater, 1997, for an exception). It is my hope that the procedure presented here, which is as simple to replicate and produces similarly impressive levels of false recognition, will take false memory research in a new direction, with more emphasis on false recognition of highly distinctive non-verbal stimuli. Further research with this procedure could help clarify the similarities and differences between true and false memories.

References

Anderson, R. (1984). Did I do it or did I only imagine doing it? *Journal of Experimental Psychology: General, 113*(4), 594-613.

Balota, D. A., Cortese, M. J., Hutchison, K. A., Neely, J. H., Nelson, D. L., Simpson, G. B., et al. (2002). *The english lexicon project: A web-based repository of descriptive and behavioral measures for 40,481 english words and nonwords.* Washington University. http://elexicon.wustl.edu/

Bartram, D. J. (1974). The role of visual and semantic codes in object naming. *Cognitive Psychology, 6*(3), 325.

Berry, C. J., Henson, R. N. A., & Shanks, D. R. (2006). On the relationship between repetition priming and recognition memory: Insights from a computational model. *Journal of Memory and Language, 55*(4), 515-533.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94*(2), 115-147.

Biederman, I., & Cooper, E. E. (1991a). Priming contour-deleted image: Evidence for intermediate representations in visual object recognition. *Cognitive Psychology, 23*(3), 393-419.

Biederman, I., & Cooper, E. E. (1991b). Evidence for complete translational and reflectional invariance in visual object priming. *Perception, 20*(5), 585-593.

Biederman, I., & Cooper, E. E. (1992). Size invariance in visual object priming. *Journal of Experimental Psychology: Human Perception and Performance, 18*(1), 121-133.

Biederman, I., & Gerhardstein, P. (1993). Recognizing depth-rotated objects: Evidence and conditions for 3-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance, 19*(6), 1162-1182.

Blaxton, T. A. (1989). Investigating dissociations among memory measures: Support for a transfer-appropriate processing framework. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 15*(4), 657-668.

Borst, G. G., & Kosslyn, S. M. (2008). Visual mental imagery and visual perception: Structural equivalence revealed by scanning processes. *Memory & Cognition, 36*(4), 849-862.

Bousfield, W. A., Esterson, J., & Whitmarsh, G. A. (1957). The effects of concomitant colored and uncolored pictorial representations on the learning of stimulus words. *Journal of Applied Psychology, 41*(3), 165-168.

Buchner, A., & Brandt, M. (2003). Further evidence for systematic reliability differences between explicit and implicit memory tests. *Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology, 56*(2), 193-209.

Cabeza, R., Burton, M. A., Kelly, S. W., & Akamatsu, S. (1997). Investigating the relation between imagery and perception: Evidence from face priming. *Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology, 50*(2), 274-289.

Cabeza, R., Rao, S. M., Wagner, A. D., Mayer, A. R., & Schacter, D. L. (2001). Can medial temporal lobe regions distinguish true from false? An event-related functional MRI study of veridical and illusory recognition memory. *Proceedings of the National Academy of Sciences of the United States of America, 98*(8), 4805-4810.

Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology, 58*(1), 17-22.

Deffenbacher, K. K., Bornstein, B. H., & Penrod, S. D. (2006). Mugshot exposure effects: Retroactive interference, mugshot commitment, source confusion, and unconscious transference. *Law and Human Behavior, 30*(3), 287-307.

DePrince, A., Allard, C. B., Oh, H., & Freyd, J. J. (2004). What's in a name for memory errors? Implications and ethical issues arising from the use of the term "false memory" for errors in memory for details. *Ethics & Behavior, 14*(3), 201-233.

Durso, F. T., & Johnson, M. K. (1979). Facilitation in naming and categorizing repeated pictures and words. *Journal of Experimental Psychology: Human Learning and Memory, 5*(5), 449-459.

Farah, M. J., Peronnet, F., Gonon, M. A., & Girard, M. H. (1988). Electrophysiological evidence for a shared representational medium for visual images and visual percepts. *Journal of Experimental Psychology: General, 117*, 248.

Fazendeiro, T., Winkielman, P., Luo, C., & Lorah, C. (2005). False recognition across meaning, language, and stimulus format: Conceptual relatedness and the feeling of familiarity. *Memory & Cognition, 33*(2), 249-260.

Feustel, T. C., Shiffrin, R. M., & Salasoo, A. (1983). Episodic and lexical contributions to the repetition effect in word identification. *Journal of Experimental Psychology: General, 112*(3), 309-346.

Finke, R. A., Johnson, M. K., & Shyi, G. C. W. (1988). Memory confusions for real and imagined completions of symmetrical visual patterns. *Memory & Cognition, 16*(2), 133-137.

Foley, M. A., Durso, F. T., Wilder, A., & Friedman, R. (1991). Developmental comparisons of explicit versus implicit imagery and reality monitoring. *Journal of Experimental Child Psychology, 51*(1), 1-13.

Foley, M. A., & Foy, J. (2008). Pictorial encoding effects and memory confusions in the Deese-Roediger-McDermott paradigm: Evidence for the activation of spontaneous imagery. *Memory, 16*(7), 712-727.

Freyd, J. J., & Gleaves, D. H. (1996). "Remembering" words not presented in lists: Relevance to the current recovered/false memory controversy. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 22*(3), 811-813.

Gallo, D. A. (2006). *Associative Illusions of Memory: False Memory Research in DRM and Related Tasks*. New York: Psychology Press.

Garry, M., Manning, C. G., Loftus, E. F., & Sherman, S. J. (1996). Imagination inflation: Imagining a childhood event inflates confidence that it occurred. *Psychonomic Bulletin & Review*, *3*(2), 208-214.

Goff, L. M., & Roediger, H. L. (1998). Imagination inflation for action events: Repeated imaginings lead to illusory recollections. *Memory & Cognition, 26*(1), 20-33.

Gonsalves, B., & Paller, K. A. (2000). Neural events that underlie remembering something that never happened. *Nature Neuroscience, 3*(12), 1316-1321.

Gonsalves, B., Reber, P. J., Gitelman, D. R., Parrish, T. B., Mesulam, M. M., & Paller, K. A. (2004). Neural evidence that vivid imagining can lead to false remembering. *Psychological Science, 15*(10), 655-660.

Hancock, T. W., Hicks, J. L., Marsh, R. L., & Ritschel, L. (2003). Measuring the activation level of critical lures in the Deese-Roediger-McDermott paradigm. *American Journal of Psychology, 116*(1), 1-14.

Heaps, C. M., & Nash, M. (2001). Comparing recollective experience in true and false autobiographical memories. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 27*(4), 920-930.

Hege, A. C. G., & Dodson, C. S. (2004). Why distinctive information reduces false memories: Evidence for both impoverished relational-encoding and distinctiveness heuristic accounts. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 30*(4), 787-795.

Henkel, L. A., & Franklin, N. (1998). Reality monitoring of physically similar and conceptually related objects. *Memory & Cognition, 26*(4), 659-673.

Henkel, L. A., Franklin, N., & Johnson, M. K. (2000). Cross-modal source monitoring confusions between perceived and imagined events. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *26*(2), 321-335.

Hicks, J., & Hancock, T. W. (2002). Backward associative strength determines source attributions given to false memories. *Psychonomic Bulletin & Review, 9*(4), 807-815.

Hicks, J. L., & Starns, J. J. (2005). False memories lack perceptual detail: Evidence from implicit word-stem completion and perceptual identification tests. *Journal of Memory and Language, 52*(3), 309-321.

Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review, 95*(4), 528-551.

Hirshman, E., Snodgrass, J. G., Mindes, J., & Feenan, K. (1990). Conceptual priming in fragment completion. *Journal of Experimental Psychology: Learning, Memory & Cognition, 16*(4), 634-647.

Hoffman, H. G. (1997) Role of memory strength in reality monitoring decisions: Evidence from source attribution biases. *Journal of Experimental Psychology: Learning, Memory & Cognition, 52*(2), 371-383.

Hyman, I. E., Husband, T. H., & Billings, F. J. (1995). False memories of childhood experiences. *Applied Cognitive Psychology*, *9*(3), 181-197.

Hyman, I. E., & Pentland, J. (1996). The role of mental imagery in the creation of false childhood memories. *Journal of Memory and Language*, *35*(2), 101-117.

Intraub, H., & Hoffman, J. E. (1992). Reading and visual memory: Remembering scenes that were never seen. *American Journal of Psychology, 105*(1), 101.

Israel, L., & Schacter, D. L. (1997). Pictorial encoding reduces false recognition of semantic associates. *Psychonomic Bulletin & Review, 4*(4), 577-581.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General, 110*(3), 306-340.

Jacoby, L. L., & Hayman, C. A. G. (1987). Specific visual transfer in word identification. *Journal of Experimental Psychology: Learning, Memory & Cognition, 13*(3), 456-463.

Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General, 118*(2), 126-135.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114*(1), 3-28.

Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review, 88*, 67-85.

Johnson, M. K., Raye, C. L., Foley, M. A., & Kim, J. K. (1982). Pictures and images: Spatial and temporal information compared. *Bulletin of the Psychonomic Society, 19*(1), 23-26.

Johnson, M. K., Raye, C. L., Wang, A. Y., & Taylor, T. H. (1979). Fact and fantasy: The roles of accuracy and variability in confusing imaginations with perceptual experiences. *Journal of Experimental Psychology: Human Learning and Memory, 5*(3), 229.

Johnson, M. K., Suengas, A. G., Foley, M. A., & Raye, C. L. (1988). Phenomenal characteristics of memories for perceived and imagined autobiographical events. *Journal of Experimental Psychology: General, 117*(4), 371-376.

Johnston, W. A., Dark, V. J., & Jacoby, L. L. (1985). Perceptual fluency and recognition judgments. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 11*(1), 3-11.

Kosslyn, S. M., & Thompson, W. L. (2000). Shared mechanisms in visual imagery and visual perception: Insights from cognitive neuroscience. In M. S. Gazzaniga (Ed.), *The New Cognitive Neurosciences* (2nd ed.). Cambridge, MA: MIT Press.

Kosslyn, S. M., Thompson, W. L., & Alpert, N. M. (1997). Neural systems shared by visual imagery and visual perception: A positron emission tomography study. *Neuroimage, 6*(4), 320-334.

Koutstaal, W., & Schacter, D. L. (1997). Gist-based false recognition of pictures in older and younger adults. *Journal of Memory and Language, 37*(4), 555-583.

Koutstaal, W., Verfaellie, M., & Schacter, D. L. (2001). Recognizing identical versus similar categorically related common objects: Further evidence for degraded gist representations in amnesia. *Neuropsychology, 15*(2), 268-289.

Lampinen, J. M., Meier, C. R., Arnal, J. D., & Leding, J. K. (2005). Compelling untruths: Content borrowing and vivid false memories. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 31*(5), 954-963.

Lampinen, J. M., Odergard, T. N., & Bullington, J. L. (2003). Qualities of memories for performed and imagined actions. *Applied Cognitive Psychology, 17*(8), 881-893.

Loftus, E. F. (1975). Leading questions and eyewitness report. *Cognitive Psychology, 7*(4), 560-572.

Loftus, E. F., Miller, D. G., & Burns, H. J. (1978). Semantic integration of verbal information into a visual memory. *Journal of Experimental Psychology: Human Learning and Memory, 4*(1), 19-31.

Loftus, E. F., & Pickrell, J. E. (1995). The formation of false memories. *Psychiatric Annals, 25*(12), 720-725.

Loftus, G. R., & Bell, S. M. (1975). Two types of information in picture memory. *Journal of Experimental Psychology: Human Learning and Memory, 104*, 103-113.

Lövdén, M., & Johansson, M. (2003). Are covert verbal responses mediating false implicit memory? *Psychonomic Bulletin & Review, 10*(3), 724-729.

Macmillan, N. A., & Creelman, C. D. (2004). *Detection Theory: A User's Guide* (2nd ed.). New York: Cambridge University Press.

Marschark, M., & Surian, L. (1989). Why are there (sometimes) concreteness effects in free recall. *Bulletin of the Psychonomic Society, 27*(6), 531-531.

Marsolek, C. J., & Andresen, D. R. (2005). Interactive visual and postvisual processes and their roles in form-specific memory. *Canadian Journal of Experimental Psychology, 59*(2), 109-123.

Mather, M., Henkel, L. A., & Johnson, M. K. (1997). Evaluating characteristics of false memories: Remember/know judgments and memory characteristics questionnaire compared. *Memory & Cognition*, *25*(6), 826-837.

McBride, D. M., Coane, J. H., & Raulerson, B. A. (2006). An investigation of false memory in perceptual implicit tasks. *Acta Psychologica*, *123*(3), 240-260.

McDermott, K. B. (1997). Priming on perceptual implicit memory tests can be achieved through presentation of associates. *Psychonomic Bulletin & Review, 4*(4), 582-586.

McDermott, K. B., Chan, J. C. K., & Weinstein, Y. (2009). *False Recognition in the DRM Paradigm on a Forced-choice Test.* Unpublished manuscript.

McDermott, K. B., & Roediger, H. L. (1994). Effects of imagery on perceptual implicit memory tests. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 20*(6), 1379-1390.

McKone, E. (2004). Distinguishing true from false memories via lexical decision as a perceptual implicit test. *Australian Journal of Psychology, 56*(1), 42-49.

McKone, E., & Murphy, B. (2000). Implicit false memory: Effects of modality and multiple study presentations on long-lived semantic priming. *Journal of Memory and Language*, *43*(1), 89-109.

Meade, M. L., Watson, M. W., Balota, D. A., & Roediger, H. L. (2006). The roles of spreading activation and retrieval mode in producing false recognition in the DRM paradigm. *Journal of Memory and Language, 56*(3), 305-320.

Metcalfe, J., & Fisher, R. P. (1986). The relation between recognition memory and classification learning. *Memory & Cognition, 14*(2), 164-173.

Michelon, P., & Koenig, O. (2002). On the relationship between visual imagery and visual perception: Evidence from priming studies. *European Journal of Cognitive Psychology, 14*(2), 161-184.

Miller, M. B., & Wolford, G. L. (1999). Theoretical commentary: The role of criterion shift in false memory. *Psychological Review, 106*(2), 398-405.

Nelson, D. L., Reed, V. S., & McEvoy, C. L. (1977). Learning to order pictures and words: A model of sensory and semantic coding. *Journal of Experimental Psychology: Human Learning and Memory, 3*, 485-497.

Nelson, D. L., Reed, V. S., & Walling, J. R. (1976). Pictorial superiority effect. *Journal of Experimental Psychology: Human Learning and Memory, 2*(5), 523-528.

Norman, K. A., & Schacter, D. L. (1997). False recognition in younger and older adults: Exploring the characteristics of illusory memories. *Memory & Cognition, 25*(6), 838-848.

Okado, Y., & Stark, C. E. L. (2005). Neural activity during encoding predicts false memories created by misinformation. *Learning & Memory, 12*(1), 3-11.

Paivio, A., Rogers, T. B., & Smythe, P. C. (1968). Why are pictures easier to recall than words? *Psychonomic Science, 11*, 137-138.

Park, S. M., & Gabrieli, J. D. E. (1995). Perceptual and nonperceptual components of implicit memory for pictures. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *21*(6), 1583-1594.

Pezdek, K., & Lam, S. (2007). What research paradigms have cognitive psychologists used to study "false memory," and what are the implications of these choices? *Consciousness and Cognition*, *16*(1), 2-17.

Potter, M. C., & Faulconer, B. A. (1975). Time to understand pictures and words. *Nature, 253*, 437-438.

Ratcliff, R., Hockley, W., & McKoon, G. (1985). Components of activation: Repetition and priming effects in lexical decision and recognition. *Journal of Experimental Psychology: General*, *114*(4), 435-450.

Reyna, V., & Brainerd, C. J. (1995). Fuzzy trace theory: Some foundational issues. *Learning and Individual Differences, 7*(2), 145-162.

Robinson, K. J., & Roediger, H. L. (1997). Associative processes in false recall and false recognition. *Psychological Science*, *8*(3), 231-237.

Roediger, H. L. (1996). Memory illusions. *Journal of Memory and Language, 35*(2), 76-100.

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 21*(4), 803-814.

Roediger, H. L., Watson, J. M., McDermott, K. B., & Gallo, D. A. (2001). Factors that determine false recall: A multiple regression analysis. *Psychonomic Bulletin & Review, 8*(3), 385-405.

Russell, W. A., & Jenkins, J. J. (1954). *The Complete Minnesota Norms for Responses to 100 Words from the Kent–Rosanoff Word Association Test.* Office of Naval Research Tech. Rep. No. 11, Contract N8 ONR 66216. Minneapolis: University of Minnesota.

Scarborough, D. L., Gerard, L., & Cortese, C. (1979). Accessing lexical memory: The transfer of word repetition effects across task and modality. *Memory & Cognition*, *7*, 3.

Schacter, D. L., Israel, L., & Racine, C. (1999). Suppressing false recognition in younger and older adults: The distinctiveness heuristic. *Journal of Memory and Language, 40*(1), 1-24.

Schacter, D. L., Verfaellie, M., & Anes, M. D. (1997). Illusory memories in amnesic patients: Conceptual and perceptual false recognition. *Neuropsychology*, *11*(3), 331-342.

Schooler, J. W., Gerhard, D., & Loftus, E. F. (1986). Qualities of the unreal. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 12*(2), 171-181.

Shepard, R. N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior, 6*, 156-163.

Shepard, R. N., & Chang, J. J. (1963). Stimulus generalization in the learning of classifications. *Journal of Experimental Psychology, 65*, 94-102.

Snodgrass, J. G., & Vanderwart, M. (1980). Standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory, 6*(2), 174-215.

Sommers, M. S., & Lewis, B. P. (1999). Who really lives next door: Creating false memories with phonological neighbors. *Journal of Memory and Language, 40*(1), 83-108.

Srinivas, K. (1993). Perceptual specificity in nonverbal priming. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *19*(3), 582-602.

Stark, C. E. L., & McClelland, J. L. (2000). Repetition priming of words, pseudowords, and nonwords. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 26*(4), 945-972.

Stern, E. R., & Rotello, C. M. (2000). Memory characteristics of recently imagined events and real events experienced previously. *American Journal of Psychiatry, 113*(4), 569.

Suzuki, K., & Takahashi, R. (1997). Effectiveness of color in picture recognition memory. *Japanese Psychological Research, 39*(1), 25-32.

Tse, C. S., & Neely, J. H. (2005). Assessing activation without source monitoring in the DRM false memory paradigm. *Journal of Memory and Language, 53*(4), 532-550.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology, 26*(1), 1-12.

Van Overschelde, J. P., Rawson, K. A., & Dunlosky, J. (2004). Category norms: An updated and expanded version of the Battig and Montague (1969) norms. *Journal of Memory and Language, 50*(3), 289-335.

Wade, K. A., Sharman, S. J., Garry, M., Memon, A., Mazzoni, G., Merckelbach, H., et al. (2007). False claims about false memory research. *Consciousness and Cognition, 16*(1), 18-28.

Warren, C., & Morton, J. (1982). The effects of priming on picture-recognition. *British Journal of Psychology, 73*, 117-129.

Weinstein, Y., & Shanks, D. R. (2008). Perceptual representations in false recognition and priming of pictures. *Memory & Cognition, 36*(8), 1415-1428.

Weldon, M. S., Roediger, H. L., Beitel, D. A., & Johnston, T. R. (1995). Perceptual and conceptual processes in implicit and explicit tests with picture fragment and word-fragment cues. *Journal of Memory and Language, 34*(2), 268-285.

Whittlesea, B. W. A. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 19*(6), 1235-1253.

Whittlesea, B. W. A. (2002). False memory and the discrepancy-attribution hypothesis: The prototype-familiarity illusion. *Journal of Experimental Psychology: General, 131*(1), 96-115.

Whittlesea, B. W. A., Jacoby, L. L., & Girard, K. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory and Language, 29*(6), 716-732.

Zeelenberg, R., & Pecher, D. (2002). False memories and lexical decision: Even twelve primes do not cause long-term semantic priming. *Acta Psychologica, 109*(3), 269-284.

Appendices

Appendix A

List of Items by Category

Animals	Clothes	Electrical Appliances	Fruit / Vegetables	Household Objects
bear [a]	*hat*	*computer*	*apple*	*drill*
cat	*shirt*	*hairdryer*	*banana*	*hammer*
dog	*shoes*	*kettle*	*broccoli*	*knife*
horse	*shorts*	*toaster*	*carrot*	*pan*
lion	*socks*	*tv*	*lettuce*	*screwdriver*
tiger	*trousers*	*washing machine*	*orange*	*spoon*
camel	bag	camera	aubergine	binoculars
cow	belt	clock	avocado	bottle
deer	bow tie	dishwasher	cherry	bowl
duck	bra	fridge	courgette	candle
eagle	coat	ipod	grapes	comb
elephant	dress	iron	kiwi	cup
flamingo	gloves	lamp	lemon	glasses
frog	sandals	laptop	nectarine	globe
giraffe	scarf	microwave	pear	hanger
kangaroo	skirt	phone	pepper	jug
peacock	slippers	photocopier	pineapple	lightbulb
rabbit	suit	printer	potato	pen
shark	tie	radio	raspberry	rolling pin
snake	t-shirt	razor	strawberry	scissors
tortoise	watch	sewing machine	sweetcorn	stapler
wolf	zip	vacuum	tomato	umbrella

[a] Italics denote critical items.

Appendix B

Baseline Bias in all Experiments

	Mean	SEM
Experiment 1a	0.37	.08
Experiment 1b	0.30	.07
Experiment 2		
Recognition	0.42	.07
Source Monitoring	0.44	.07
Experiment 3		
Perceptual	0.40	.06
Conceptual	0.27	.06
Experiment 4	[0]	–
Experiment 5	-0.01	0.10
Experiment 6	0.25	0.05
Experiment 7	-0.00	0.07

Appendix C

Comparison of Direct and Indirect Tests in Experiments 5-7

	Direct Test: Proportion called "old"		Indirect Test: Priming effect (msec)	
Experiment	True Memory[a]	False Memory[a]	True Memory[b]	False Memory[b]
Experiment 5	0.71	0.07	136	10
Experiment 6	0.74	0.26	75	-37
Experiment 7	0.76	0.48	123	-6

[a]On the direct test, true memory represents hits to seen items minus false alarms to new items, while false memory represents false alarms to imagined (or word) items minus false alarms to new items.

[b]On the indirect test, true memory represents priming of seen items compared to the new item baseline, while false memory represents priming of imagined (or word) items compared to the new item baseline.

Footnotes

[1]Bias was calculated from hit and false alarm rates rather than the raw confidence data due to the low number of high confidence false alarms to new items.

[2]This assumption receives further support from the fact that participants claimed to guess (i.e., indicated a confidence level of 1, *not at all sure*) on .10 of all seen-new trials and .11 of all seen-imagined trials, both of which were not significantly different to .1, or 1 out of 10 trials (both ps $> .5$).

[3]The same analysis was also carried out on the subset of word and new items to which participants made false alarms. In order to obtain a reliable measure of RT, only the nine participants for whom at least 4 false alarm RT observations were available in both the word and new item conditions were included in the analysis. Participants were not reliably faster at identifying word items (M = 1173 msec) than new items (M = 1196 msec), $t < .5$.

[4]It was not possible to look at the subset of studied word and new items which were accompanied by an old response due to the low rate of false alarms to new items. Instead, RTs of studied word items were compared to the new item RT baseline. The 20 participants for whom at least 4 false alarm RT observations were available in the studied word condition were included in the analysis. The effect was non-significant and in the opposite direction to that predicted, with studied word items (M = 1552) identified slower than new items (M = 1482), $t(19) = -1.60$, $p = .13$.

[5]As for Experiment 5 false alarms to new items were rare so RTs of studied word items were compared to the new item RT baseline. The 18 participants for whom at least 4 false alarm RT observations were available in the studied word condition were not reliably faster at identifying those items (M = 1401 msec) than new items (M = 1420 msec), $t < 0.5$.