Memory blindness: Altered memory reports lead to distortion in eyewitness memory

Kevin J. Cochran¹ · Rachel L. Greenspan¹ · Daniel F. Bogart¹ · Elizabeth F. Loftus¹

Published online: 16 February 2016
© Psychonomic Society, Inc. 2016

Abstract Choice blindness refers to the finding that people can often be misled about their own self-reported choices. However, little research has investigated the more long-term effects of choice blindness. We examined whether people would detect alterations to their own memory reports, and whether such alterations could influence participants' memories. Participants viewed slideshows depicting crimes, and then either reported their memories for episodic details of the event (Exp. 1) or identified a suspect from a lineup (Exp. 2). Then we exposed participants to manipulated versions of their memory reports, and later tested their memories a second time. The results indicated that the majority of participants failed to detect the misinformation, and that exposing witnesses to misleading versions of their own memory reports caused their memories to change to be consistent with those reports. These experiments have implications for eyewitness memory.

Keywords False memory · Choice blindness

When asked about the motives for their behavior, reasons for their choices, and sources of their memories, people can often produce explanations. But are these explanations the true origins, or are they *post-hoc* constructions based on plausible inferences made from the available evidence (Bem, 1972)? A growing body of research on “choice blindness” suggests that peoples' introspective abilities can be quite limited. The

literature shows that when people are asked to choose between several options, they often fail to notice if they are then given one of the nonchosen options. In the present research, we extended the choice blindness finding to the novel domain of eyewitness memory. More specifically, we asked whether people are able to detect changes made to their own previously given memory reports, and whether such changes affect what people subsequently remember.

Choice blindness

In the choice blindness paradigm, people are first given a choice between several options—for instance, they might be asked to taste two types of jam and indicate which they prefer (Hall, Johansson, Tärning, Sikström, & Deutgen, 2010). Next, they are given the option that they picked and asked to explain why they made that choice. Unbeknownst to participants, due to a concealed manipulation, the option they are given is not the choice they had originally selected—rather, it is one of the nonchosen options. For instance, a participant who had initially indicated that grapefruit jam was his favorite might be presented with the cinnamon apple jam as if it were the one he had chosen (and vice versa). The finding of interest is that people often fail to notice this manipulation, come to endorse the option they had initially rejected, and even confabulate reasons why they made a choice that they never really made. When given a different flavor of jam from the one they had truly favored, only one third of participants displayed any evidence of having noticed.

High rates of blindness are by no means limited to superficial choices. Although some variables, like the type of decision, the manner in which the manipulation occurs, and other between-experiment variables, seem to produce different rates of blindness, choice blindness has been shown to be robust

✉ Kevin J. Cochran
kjcochra@uci.edu

¹ Department of Psychology and Social Behavior, University of California, Irvine, CA 92697-7085, USA

across a variety of domains. Participants' blindness rates have been remarkably high in studies of their political and moral attitudes (53 %; Hall, Johansson, & Strandberg, 2012), financial decision making (63 %; McLaughlin & Somerville, 2013), and even their reported histories of criminal and norm-violating behavior (8 %–10 %; Sauerland et al., 2013).

Often in choice blindness studies, the dependent variable of interest is simply the proportion of participants who fail to notice the manipulation. However, several recent studies have investigated whether choice blindness might have lasting effects (Johansson, Hall, Tärning, Sikström, & Chater, 2014; Merckelbach, Jelicic, & Pieters, 2011). The theoretical importance of this question is immense: The choice blindness literature has shown that people can often be misled in the short term about their own preferences, but these more recent studies demonstrate that this simple manipulation can often have lasting consequences on peoples' attitudes and behaviors.

This recent trend in the literature raises an important question: Can choice blindness have lasting effects on eyewitness memory in the same way it influences peoples' attitudes? That is, if people are misled about their own previous memory reports, will that manipulation affect their future memories? To our knowledge, only one published study has examined the subsequent effect that choice blindness has on eyewitness memory (Sagana, Sauerland, & Merckelbach, 2014). The researchers found that participants were concurrently blind to changes in their suspect identifications between 33 % and 68 % of the time, depending in part on the amount of time between when participants first made an identification and when they were misled about that identification (see below). However, in one experiment, the researchers found almost no memory distortion. In another experiment, they unfortunately obtained a low response rate to their follow-up questionnaire, and consequently their test of whether choice blindness had influenced participants' memories was limited. The researchers concluded that "future studies on choice blindness and eyewitness identification might profit from an explicit consideration of the misinformation literature" (Sagana et al., 2014, p. 762).

The misinformation effect

The misinformation effect is the finding that if people receive misleading information about a previously witnessed event, they will often incorporate that misinformation into their memories of the event (Loftus, Miller, & Burns, 1978; for a review, see Loftus, 2005). In one study, when participants viewed a depiction of a car traveling through an intersection with a stop sign, if they were later exposed to a suggestive question that mentioned a yield sign, many falsely remembered a yield sign in a subsequent forced choice recognition test (Loftus et al., 1978). In most misinformation studies, the

misinformation is either presented in a similar surreptitious way or attributed to a third party, such as a putative "co-witness" who is actually a confederate of the experimenters (Meade & Roediger, 2002). Some studies have employed participants themselves in aiding with inducing memory distortion (e.g., Roediger, Jacoby, & McDermott, 1996), but it remains an open question whether we can elicit a misinformation effect from participants by misleading them about their own previous memory reports.

In many choice blindness studies, the manipulation occurs immediately after a decision has been made, which is advantageous because it makes for a compelling effect and participant experience. By using sleight-of-hand tricks and rigged props, researchers can make participants' responses change practically "before their eyes" (e.g., Hall et al., 2012; Hall et al., 2010). But research on the misinformation effect has shown that the durations of the intervals between the original event and the manipulation and between the manipulation and the test can have important consequences for the strength of the misinformation effect. Misinformation has the greatest influence on participants' memories when it is presented after a long retention interval and immediately prior to a memory test, since the original memory has decayed, whereas the misinformation is fresh (Loftus et al., 1978). A parallel effect could be found for choice blindness; after a longer delay, participants might have more ambiguous memories for their choices, and when asked to justify their choices, they may rely more on a constructive process of evaluating plausible reasons why they may have made a given choice (Bem, 1972; Johansson, Hall, Sikström, Tärning, & Lind, 2006; Sagana et al., 2014). Thus, for the present studies, we included a longer retention interval between the original event and the presentation of misinformation.

Another factor in misinformation research linked to the timing of the manipulation is whether participants detect that the postevent information they receive is inaccurate (Loftus, 2005; Tousignant, Hall, & Loftus, 1986). According to the discrepancy detection principle, peoples' memories for events are more likely to change if they fail to notice this discrepancy. By using a longer interval between when participants make a choice and when they are exposed to the misinformation, we should be able to limit participants' ability to notice the discrepancy.

Choice blindness and the misinformation effect share many characteristics. In the misinformation paradigm, participants first witness an event, are then exposed to misleading information, and are finally tested on their memory for the event. In choice blindness, participants first make a decision, express an attitude, or choose an option; are then exposed to false feedback about their choice; and finally are tested implicitly on their acceptance of this manipulation. One major difference between these two paradigms is in the agency of the subject. Traditionally in misinformation studies, participants are more

passive consumers of information; they witness events and receive misleading information. In choice blindness studies, by contrast, participants express their preferences or make decisions between options before receiving misleading information *about those preferences or decisions*. The present experiments can help to shed light on whether participants will still exhibit the misinformation effect when misinformed not about information they have passively consumed, but rather about decisions they actively made and memories they actively reported.

There is some reason to believe that misleading people about their own memory reports may produce a diminished misinformation effect. Memory is subject to social influences, such as pressures to conformity and informational social influence; indeed, when witnesses are allowed to discuss their memories of an event with each other, they sometimes exhibit a “memory conformity” effect, in which their initially disparate memories become more alike (Gabbert, Wright, Memon, Skagerberg, & Jamieson, 2012). Participants who are told how another putative participant responded may succumb to these social influences, but participants who are misled about their own previous responses should not. In other words, misinformation about one’s own memory report is void of any social information, and thus does not exert social pressure that may influence memory. Therefore, it is important to test whether this “self-sourced” misinformation causes memory distortion at all, and if it does, to compare the effects of “self-sourced” misinformation with misinformation attributed to another witness.

In the present experiments, we sought to integrate choice blindness and the misinformation effect. In Experiment 1, participants first witnessed an event and were then asked questions testing their memories for episodic details of the event. Later, they were shown their own memory reports, but some of their responses had been altered. Finally, they were asked the memory questions a second time, in order to determine whether the misinformation had caused memory distortion. In Experiment 2 we followed a similar procedure, but participants’ memory task was to identify the suspect out of a photo lineup. After receiving misinformation about their selection, participants were asked a second time to select the suspect from a lineup.

Experiment 1

Method

Participants A group of 186 students at a large university in southern California participated in exchange for partial course credit. Six of the participants failed to complete the experiment, and 15 failed an attention check, yielding a final sample of 165. The sample size was determined by previous

experience with research on the misinformation effect; no data collection stopping rule was in place.

Design Experiment 1 consisted of two experimental conditions in which participants received falsified versions of their own memory reports. In the “self-sourced” condition, these reports were presented as though they were the exact account that the participant him- or herself had previously given. In the “other-sourced” condition, these reports were presented as though they were the accounts another participant had reported in a previous trial.

The memory reports consisted of ten items. For each participant, three of those items, chosen at random, were manipulated (misinformation items), whereas the other seven items were not manipulated (control items). Thus, the present study was based on a 2 (self-sourced vs. other-sourced, between participants) $\times$ 2 (misinformation vs. control, within participants) mixed design. The dependent variable of interest was how much participants’ reports of their memories would change between the baseline test and the final test in the direction of the misinformation.

Procedure The present study was conducted online. Participants watched a short slideshow adapted from Okado and Stark (2005), depicting a female character interacting with three other characters, one of whom steals her wallet. The participants then completed personality measures during a retention interval of approximately 15 min, which contributed to the credibility of our cover story. Next, participants were asked about their memories for the slideshow (Test 1). Each participant was asked the same ten questions, displayed one at a time, in a random order. The questions were designed to simulate those that police might have for real eyewitnesses, such as “what color was the thief’s jacket,” or “how tall was the thief,” with responses ranging from shades of green to shades of blue for the former, and from five feet seven inches to six feet two inches for the latter. The ten memory questions were all presented on 15-point Likert-type scales.

After a second 15-min retention interval, participants entered the misinformation stage. They were shown their responses to the memory questions, but three of their responses, chosen randomly, had been altered. For these three items, the participants’ answers were shifted by four points along the Likert scales. The direction that each response was shifted was randomized, unless the initial responses were too close to the endpoints of the scale to allow for a shift of four points in one direction, in which case responses were shifted toward the center. For each critical question, a difficult or impossible follow-up was developed; participants were shown their previous responses, presented either as their own reports or as another participant’s reports, and then asked the follow-up question. For example, one page read “In a previous trial, another participant said that the thief’s jacket was the color

indicated,” with an arrow pointing to one of the 15 color swatches. When this was a control item, the arrow pointed to the color swatch that the participant had selected earlier, but when this was a misinformation item, the arrow pointed to a color swatch four spaces over. On the next page they were asked “What brand was it?” In this way, participants were required to engage with the misinformation. The ten items were randomized, and the misinformation appeared for the fourth, sixth, and ninth items.

After a final 15-min retention interval, the final stage of the experiment involved participants responding to the same ten memory questions a second time. The questions were displayed in the same order they had been in on the misinformation stage. At the end of the study, participants were debriefed. First they were asked what they thought the study was about in a multiple-choice question with four options: “how your personality affects your visual perception” (the cover story for the experiment), “how your personality affects your memory,” “how misleading information affects your memory,” and “the difference between short-term and long-term memory.” Participants were then asked whether anything in the experiment had seemed odd to them, and they were given room to explain their answer.

Data analysis To determine whether participants were influenced by the misinformation, we analyzed the mean differences in participants’ responses to the misinformation items versus the control items. When participants changed their responses in ways congruent with the misinformation, they received positive scores, and when their responses changed away from the misinformation, they received negative scores. This type of analysis, used by Merckelbach et al. (2011), is advantageous in that it is sensitive to both the magnitude and the direction of the change in memory (i.e., consistent or inconsistent with the misinformation).

Results

Blindness to hypotheses When asked in a four-option multiple-choice question what they thought the experiment was about, 24 % of participants selected the true purpose of the experiment, “how misleading information affects your memory,” whereas 40 % of participants selected the response associated with our cover story, “how your personality affects your visual perception,” with the remaining participants choosing one of the foil options, “how your personality affects your memory” (31 %) or “the difference between short-term and long-term memory” (5 %). Because the proportion of participants who correctly identified the purpose of the experiment was so close to what we would expect by chance alone, it is unclear whether these participants were truly aware of the study’s purpose or were simply guessing. When asked whether anything in the experiment struck participants as odd, only

18 % of the participants reported finding anything odd, and only seven participants (4 % of the sample) mentioned anything specifically related to the purpose of the study.

The broadest measure of the number of participants who detected the purpose of the study includes both those who guessed the purpose in the multiple-choice question and those who reported finding something odd about the study. This measure almost certainly overestimates the true number of detectors, but it is useful nevertheless to bound the estimates of participant detection. By this measure, 60 participants, or 36 % of the sample, were suspicious of the purpose of the study on some level. By contrast, the narrowest measure of detectors is the percentage of participants who specifically mentioned something related to the hypotheses of the study when asked whether anything about the study struck them as odd. By this measure, seven participants, or 4 % of the sample, detected the purpose of the study. The percentage of participants who truly detected the hypotheses of the study is likely between these two extremes. The seven participants who identified something specifically related to the hypotheses of the study were included in the below analyses. The general pattern of results remained the same whether or not these participants were included in the analyses.

Change in memory In this analysis, all misinformation items were treated as though participants had been misled positively—that is, to the right on the scale. Trials on which participants were misled negatively were reverse coded.

The change in memory reports in the predicted direction is shown in Fig. 1. As is shown in the figure, for the control items, participants’ responses did not appear to change from Test 1 to Test 2. For misinformation items, however, the responses did appear to change between the two tests, and the magnitudes of these changes were similar for the self-sourced

Fig. 1 Mean scores for memory items. Scores represent the average responses that participants gave for memory items. For this analysis, all responses were coded as though participants were misled toward the positive (right) side of the scale. Changes in mean scores can be calculated by subtracting the Time 1 scores from the Time 2 scores for each condition. Error bars represent ± 1 SEM

and the other-sourced groups. To analyze these results, a 2 (Source: self vs. other) $\times$ 2 (Item Type: misinformation vs. control) $\times$ 2 (Time: Test 1 vs. Test 2) repeated measures analysis of variance (ANOVA) was used. These analyses revealed a significant main effect for misinformation items, $F(1, 163) = 9.89, p = .002, \eta_p^2 = .06$, 90 % CI for effect size = [.01, .12]. We also found a significant effect of time, $F(1, 163) = 94.61, p < .001, \eta_p^2 = .37$, 90 % CI for effect size = [.27, .45]. Finally, there was a significant Time $\times$ Misinformation interaction, $F(1, 163) = 78.88, p < .001, \eta_p^2 = .33$, 90 % CI for effect size = [.23, .41]. For misinformation items but not control items, participants' memories at Time 2 were shifted in the direction of the misinformation. No effects were found for misinformation source, all $ps > .05$.

Discussion

Experiment 1 demonstrated that when witnesses were exposed to altered versions of their own memory reports for episodic details of an event, their memories changed to be consistent with those altered reports. Manipulated items produced a greater change in memory than did control items for both the “self-sourced” misinformation group and the “other-sourced” group, and this change was consistent with the misinformation the participants received.

One interesting question that these results raise is whether those who detected the manipulation exhibited a weaker misinformation effect. According to the discrepancy detection principle (Tousignant et al., 1986), this should be the case: Those participants who noticed the discrepancy between their initial report and the misinformation should be less likely to experience memory distortion. Although only 4 % of the sample explicitly indicated knowledge of the hypotheses of the study when asked if they had found anything strange in the experiment, 24 % selected the correct response from a multiple-choice question asking participants what they thought the true purpose of the study was. Unfortunately, because of the nature of our questions, it was difficult to examine this possibility; the question asking participants whether they found anything odd was optional—that is, participants were not forced to respond. Only 18 % of the participants responded at all, and it is unclear whether the other 82 % truly did not find anything odd or were simply trying to complete the study more quickly. Participants did have to answer the multiple-choice question asking them what they thought the experiment was about, but it is unclear whether the 24 % of participants who selected the correct answer truly understood the nature of the experiment or were simply guessing; by chance alone, 25 % of participants would be categorized as “detectors.” We think these measures are useful for bounding our estimates of detection—the true proportion of participants who detected the discrepancy is likely between our two extreme measures—but they have questionable utility beyond that.

Given the limitations of our measure of detection in Experiment 1, one important addition to Experiment 2 was the inclusion of a more precise measure of concurrent detection, which allowed us to examine whether any of the observed effects were due to exposure to the choice blindness manipulation per se, or instead were due to participants' failure to detect such a manipulation (Sagana et al., 2014). Additionally, in Experiment 2 we sought to extend the findings of Experiment 1 by using a different memory task; rather than testing their memories for the episodic details of a witnessed event, in Experiment 2 we tested participants' abilities to identify a suspect from a lineup. Finally, Experiment 2 was designed to be fully between participants, which allowed us to avoid cascade effects, by which participants who detected one manipulation could scrutinize subsequent trials (Johansson, Hall, Sikström, & Olsson, 2005).

Experiment 2

Method

Participants A total of 392 students at a large university in southern California participated in exchange for partial course credit. Due to technical issues, 13 of the participants did not watch the critical slideshow and were excluded from the analysis, leaving a final sample of 379. Using previous experience with misinformation research, we collected a large enough sample to ensure that analyses could be conducted on important subgroups. We planned to stop data collection after collecting between 350 and 400 valid responses.

Design The present study had three conditions: control, confirming information (called “nonmanipulated” in choice blindness studies), and manipulated. In the *control* condition, participants received no feedback about their identification. In the *confirming information* condition, they received accurate feedback about their identification. In the *manipulated* condition, they received misleading feedback about their identification decision. The misleading feedback was presented as if it were the participant's own prior identification.

The slideshow used in this study depicted a Caucasian man stealing a radio from a car. The man's face was in view for 18 s. Lineup photographs were taken from two databases: the Psychological Image Collection at Stirling (<http://pics.stir.ac.uk>) and the Center for Vital Longevity Face Database (Ebner, 2008). All of the photographs were in color on a white background. The faces were pilot tested in order to create a lineup of relatively dissimilar faces, so that changes in identification decision from Lineup 1 to Lineup 2 could be attributed to the manipulation rather than to confusion due to facial similarity. A group of 24 participants rated the similarity of pairs of faces on a 7-point scale

(1, *not at all similar*; 7, *highly similar*). The mean similarity of the final sample of faces ranged from 1.87 to 3.41 ($M = 2.53$, $SD = 0.77$).

Procedure The present study followed a procedure similar to that of Experiment 1. The study was conducted online. Participants first watched the slideshow and then completed memory tasks consistent with the cover story during a retention interval of about 10 min. Next, participants viewed Lineup 1, which was a six-person, target-absent lineup. Photographs were presented in random order in two columns, and participants were not given the option to reject either lineup. Example photographs are shown in Fig. 2. Following their lineup decision, participants rated their confidence on an 11-point scale (1 = 0 % *confident my decision was correct*, 11 = 100 % *confident my decision was correct*).

Participants then completed another 10-min retention interval, followed by the critical manipulation. At this point, participants were randomly assigned to one of the three conditions: control, confirming information, or manipulated. Those assigned to the confirming information condition read the following statement: “Earlier in the study, you picked the photo of the man you saw in the slideshow. On the next page, you will briefly see the photo of this person.” When participants advanced to the next page, the photograph they picked was shown for 4 s. After this, participants were presented with a free-response question in which they were asked to explain why they had selected that person from the lineup.

Participants in the manipulated condition viewed the same instructions as those in the confirming information condition. They were told that they would see a photograph of the man they had selected from the lineup. However, when they advanced to the next page, the photograph shown was a randomly selected, nonchosen option from Lineup 1. After this, participants received the same instructions as in the confirming information condition.

In the control condition, participants were not shown a photograph. Prior to the free-response question, control participants were asked to think back to when they had selected the man they saw from the slideshow. They were then asked to explain why they had picked this person. The instructions for this task were the same as the instructions given after the

photograph was shown in the confirming information and manipulated conditions.

Following another retention interval similar to the previous ones, participants completed Lineup 2. Lineup 2 was identical to Lineup 1, except that the order of faces was randomized. Confidence in this choice was assessed with the same scale described previously. Finally, participants completed a basic demographics questionnaire and a funneled debriefing to assess retrospective detection. This funneled debriefing began with broad, open-ended questions about the study, and followed up with increasingly more specific, multiple-choice questions about whether the participants realized the true purpose of the study. As in Experiment 1, participants were asked about the true purpose of the study and whether they had found anything odd about the experiment. For those not in the control condition, two further questions were asked. First, they were asked whether they would notice a switch in the photographs if this were done in a similar experiment (see Johansson, Hall, & Sikström, 2008). Finally, they were told about the manipulation and asked whether they had noticed it at the time.

Results

Detection Concurrent detection was measured by coding the participants’ free responses when they described their reasons for making their identification. Two research assistants, blind to the hypotheses, were trained in evaluating participants’ responses, to assess whether participants had detected the manipulation (see Table 1 for examples). The two raters agreed on 99.5 % of the responses; disagreements were resolved by a third rater. Zero responses in the confirming information and control conditions were rated as “concurrently detected,” providing further indication of the accuracy of the coding. In the manipulated condition, 47.2 % of participants had concurrently detected the manipulation.

Detectors and nondetectors differed in the numbers of words they wrote during the free-response section, $t(125) = 2.73$, $p = .007$. Detectors ($M = 23.12$, $SD = 21.75$) wrote fewer words than nondetectors ($M = 33.48$, $SD = 21.02$). This effect was likely due to the fact that a majority of detectors simply responded that the picture they were shown was not the one they chose and did not elaborate further. A one-way ANOVA with number of words written serving as the dependent variable, and condition separated by blindness (control, confirming information, detectors, and nondetectors) serving as the independent variable, revealed significant differences between the groups, $F(3, 378) = 9.25$, $p < .001$. Post-hoc analyses revealed that detectors wrote significantly fewer words than participants in the control and confirming information condition, $ps < .001$. However, nondetectors did not significantly differ from the participants in the control and confirming information conditions, $ps > .10$. Nondetectors

Fig. 2 Example lineup members

Table 1 Example statements of detectors and nondetectors

Detectors	Nondetectors
This was not the picture I selected.	I picked this photo because the guy looked very mono-toned. He had a sort of wide headed head and face.
i dont recall choosing this person, i believe it was someone else	I think he has the same facial construction as the guy in the slideshow.
I didn't even select this photo—I chose the guy with the only decent looking hair. . . . I see what you're trying to do.	I couldn't decide which picture resembled the guy in the slideshow the most, but I went with this one. His sideburns seem to relate the original, and his eyes were not too dissimilar.
I don't think that was the person I picked out. However, If that really was the person I picked out then I guess there are similar features like the jawline and the nose. The hair threw me off a bit.	I selected this picture because he looked like the man presented in the slideshow. The men in the other pictures either looked too thin or too heavy.
I did not pick this guy . . . so I am a bit confused as to why it said I picked him. I was debating on him and the guy I actually chose, and I went with the other guy because of having read what it said that it may not look exactly like the original.	I chose the picture because based on my memory I thought it was the person from the video. I thought his hair style and face shape was similar to the man in the video.

The original spellings are included

wrote the same number of words when describing their reasons for their identification as participants in the control and confirming information conditions, despite having written about a nonchosen target.

Memory change from Lineup 1 to Lineup 2 Memory change was operationalized through participants' consistency at Lineup 1 and Lineup 2. Memory change occurred when participants selected a different lineup member for Lineup 2 than they had selected at Lineup 1. Participants who made the same identification for Lineups 1 and 2 were coded as showing no memory change.

Overall, 25 % of participants showed evidence of memory change. The rates of memory change differed by condition: 17 % for the control condition, 22.2 % for the confirming information condition, and 34.6 % for the manipulated condition. A logistic regression was run to determine whether these differences were significant. Lineup change (change or no change) served as the dependent variable, and condition (control, confirming information, and manipulated) served as the sole categorical independent variable. The confirming information condition served as the reference group. The overall model was significant, $\chi^2(2, N = 379) = 10.51, p = .005$. Change rates did not differ significantly between the control and confirming information conditions, $OR = 0.74, p = .344$, 95 % CI [.397, 1.38]. Participants in the manipulated condition had a significantly higher rate of changing than did those in the confirming information condition, which suggests that presenting participants with misinformation caused significant memory change, $OR = 1.86, p = .03$, 95 % CI [1.06, 3.24].

However, this analysis may be misleading, since it aggregates participants in the manipulated condition who detected the manipulation with those who were blind to the manipulation (Johansson et al., 2014). For detectors, only 13.3 % changed from Lineup 1 to Lineup 2. This result differs significantly from the nondetectors, who changed 53.7 % of the time, $\chi^2(1, N = 127) = 22.82, p < .001, \phi^2 = .42$. A second logistic regression was run to investigate whether detectors and nondetectors differed from the other two conditions. Since the first regression had revealed no significant differences between the control and confirming information conditions, these conditions were collapsed. Lineup change served as the dependent variable, and group (control/confirming information, detectors, and nondetectors) served as the sole categorical independent variable. Control/confirming information served as the reference group. The overall model was significant, $\chi^2(2, N = 379) = 33.88, p < .001$. There was no significant difference between detectors and the control/confirming information groups, $OR = 0.62, p = .248$, 95 % CI [0.28, 1.39]. Nondetectors switched their identification from Lineup 1 to Lineup 2 significantly more than did participants in the control/confirming information groups, $OR = 4.69, p < .001$, 95 % CI [2.65, 8.31].

When participants in the manipulated condition changed their identification at Lineup 2, they mostly changed in the direction of the manipulation. That is, of those in the manipulated condition who switched, 57 % switched to the face implicated by the misinformation. This rate of change to the target face is significantly greater than we would expect by chance, $\chi^2(1, N = 94) = 13.58, p < .001$.

Discussion

Experiment 2 replicated the findings of Experiment 1 and extended those results by measuring concurrent detection. This allowed for a more direct test of the discrepancy detection principle. The results supported the principle—nondetectors were significantly more likely to show evidence of memory change than were detectors. This indicates that blindness to the manipulation, and not the mere presentation of the misinformation, is what really drives memory distortion (Sagana et al., 2014). When participants detected the misinformation, their responses were similar to those when the misinformation was not presented at all.

In the present study, we used only target-absent lineups. Although such lineups may not always approximate real-world lineups, the use of target-present lineups in the present study would have added a confound, making the results difficult to interpret. If a target-present lineup were used, then participants who initially made a correct identification would have received misinformation that led them away from the correct answer, whereas the remaining participants would have received misinformation that led them from one foil target to another. The use of a target-absent lineup ensured that all participants received the same experimental treatment. The lineup instructions used in the present study did not specify that the suspect might or might not be in the lineup. One way to approximate those participants who would have chosen to reject the lineup would be to examine the participants' confidence in their identifications: Participants who reported low confidence in their initial identification might have chosen to reject the lineup if they were given the option. Excluding participants at the lowest two levels of confidence did not change the detection rate (45 %), and all results that were previously statistically significant remained significant, $ps < .05$.

One might be tempted to argue that nondetectors were simply not attending to the materials as attentively as the detectors. Our data do not support this conclusion. If nondetectors were simply not paying attention, then we would expect their identifications on Lineup 2 to be somewhat random. This was not the case. The majority of nondetectors at Lineup 2 chose the face implicated by the misinformation, suggesting that they were incorporating the misinformation into their memory for the event.

General discussion

In the present studies, we sought to address the question of whether an eyewitness could develop false memories for an event by being exposed to a fabricated version of his own memory report. In doing so, we integrated the phenomena of choice blindness and the misinformation effect. Experiment 1

demonstrated that the misinformation effect could be elicited from participants by telling them they had reported remembering episodic details in a different way from how they had earlier reported remembering those details. Experiment 2 generalized these findings to another memory task, eyewitness identification, and demonstrated that blindness to the manipulation, rather than mere exposure to the manipulation, drives subsequent memory change, consistent with previous theoretical and experimental work (Johansson et al., 2014; Sagana et al., 2014; Tousignant et al., 1986). We call this novel consequence of choice blindness on eyewitness memory *memory blindness*: When witnesses are exposed to manipulated versions of their own memory reports, they often fail to notice the manipulation, and their memories often change to be consistent with those altered reports.

The present experiments demonstrated the long-term effects that choice blindness might have for eyewitness memory, and this memory blindness effect could have important practical implications for the legal system. In criminal investigations, witnesses are sometimes handed summaries of their statements and asked to sign them. But if those summaries contain errors, whether due to clerical errors or deliberate manipulation, by merely reviewing their own statements, witnesses might contaminate their memories. Although lay people may believe that they would notice a discrepancy between what they reported and the content of their altered statement, the present findings suggest that many witnesses may fail to notice, and that such a failure can cause their memories to change to be consistent with those altered statements.

Some readers will notice a similarity between our paradigm and retrieval-enhanced suggestibility (RES; Chan, Thomas, & Bulevich, 2009). Briefly, in RES studies, some participants are asked to recall information about an event twice: once immediately after witnessing it, and once after encountering postevent information (PEI) regarding the event. Relative to participants who are only tested once (i.e., after encountering PEI, but not before), participants who are tested twice display greater levels of suggestibility to misinformation, perhaps because initially recalling the event makes the memory traces more susceptible to distortion. However, RES studies use traditional misinformation procedures; participants are exposed to misinformation via an audio narrative summarizing the event (Chan et al., 2009). In our experiments, we used a choice blindness paradigm in which participants were misinformed about their own previous memory reports. It would be an interesting study indeed that disentangled the RES effect from the memory blindness effect that we observed. Unfortunately, such a disentangling is beyond the scope of our experiments, so we leave it to future research.

Given the existing literature on choice blindness and the decades of research on the malleability of memory, are our findings really that surprising? Johansson et al. (2014) found that when participants selected which of two faces they found

more attractive, if they were then exposed to the choice blindness manipulation, many would later find the initially unselected face to be more attractive. In other words, the manipulation caused a change in preference. But this is not the same as a change in memory. Some explanations for this type of finding, including cognitive dissonance, compellingly explain demonstrations of attitude change, but they may require additional steps to apply to memory change. Unlike attitudes, memories reflect (at least partially) some ground truth—they are representations (though sometimes distorted ones) of events that actually happened. Attitudes are more subjective and less constrained. Other researchers have discussed self-perception theory as an explanation for similar findings (Pärnamets, Hall, & Johansson, 2015), and we agree that self-perception theory could explain our findings, as well. But the present experiments were based on different memory tasks in an eyewitness context. In light of the unique differences in the present experiments, we submit that there is no way that one could have known a priori that the memory blindness manipulation would be effective at altering people's memories, as the choice blindness manipulation is with altering preferences.

One limitation to the present experiments was that they were conducted entirely online. Participants may have been paying less attention to the study materials than they would in the lab, or they may have failed to follow our instructions. However, we tried to eliminate this possibility by including attention checks in our studies, which allowed us to exclude some participants for not attending to the study materials. A second issue with online data collection is that it makes for a less compelling choice blindness manipulation than when a survey inexplicably changes responses (Hall et al., 2012) or when one flavor of jam is magically replaced by another (Hall et al., 2010). Nevertheless, we obtained relatively low rates of detection of our manipulations, which suggests that online data collection is a valid way to study choice blindness (see also Johansson, Hall, Gulz, Haake, & Watanabe, 2007). Online data collection also confers some important benefits for studying choice blindness. Participants might be unwilling to report detecting a manipulation in person, either for fear of appearing foolish, for fear of ruining the experiment, or because they do not want to create trouble for the researcher. Online, these worries might be mitigated.

Another limitation to the present study concerns how we measured detection. In Experiment 1, we only measured retrospective detection, and although we followed a “funneled debriefing” procedure, our questions may not have truly assessed detection. Thus, for Experiment 2, we coded participants' responses to an open-ended question for evidence of concurrent detection. But even this measure, used ubiquitously in choice blindness research, has its limitations (Sagana et al., 2014). Participants might misunderstand the instructions, they might respond carelessly, or they might detect the

discrepancy but fail to report it. Future research on choice blindness should investigate other, perhaps more implicit, methods of measuring detection (Fazio & Olson, 2003). For instance, Johansson et al. (2006) used word frequency and latent semantic analysis to examine potential differences in the language that participants use to justify their choices for manipulated versus nonmanipulated trials. Future studies might examine whether participants require more time to frame an explanation for a manipulated than for a nonmanipulated choice, since, hypothetically, justifying a manipulated choice should require more effort. Another possibility would be to examine participants' facial expressions or physiological reactions during a manipulated trial. Some participants might feel as though something was not quite right in the experiment, but this feeling might not be specific enough or motivating enough for them to report detecting the manipulation. If, as we discussed above, peoples' introspective abilities are indeed quite limited, then examining processes that occur outside of awareness might provide a fruitful way for measuring detection, and might ultimately lead to more valid and more precise measures of blindness.

Author note For Experiment 1, K.J.C., D.F.B., and E.F.L. developed the study concept and contributed to the study design. Testing and data collection were performed by K.J.C., and both K.J.C. and D.F.B. performed the data analysis and interpretation, with input from E.F.L. For Experiment 2, R.L.G., K.J.C., and E.F.L. developed the study concept and contributed to the study design. Testing and data collection were performed by R.L.G., and both R.L.G. and K.J.C. performed the data analysis and interpretation, with input from E.F.L. K.J.C. and R.L.G. drafted the manuscript, and all authors provided critical revisions. All authors approved the final version of the manuscript for submission. Support for this research was provided by a Fellowship from the Center for Psychology & Law at the University of California, Irvine. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1321846.

References

- Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 6). New York, NY: Academic Press.
- Chan, J. C. K., Thomas, A. K., & Bulevich, J. B. (2009). Recalling a witnessed event increases eyewitness suggestibility: The reversed testing effect. *Psychological Science*, 20, 66–73. doi:10.1111/j.1467-9280.2008.02245.x
- Ebner, N. C. (2008). Age of face matters: Age-group differences in ratings of young and old faces. *Behavior Research Methods*, 40, 130–136. doi:10.3758/BRM.40.1.130
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology*, 54, 297–327. doi:10.1146/annurev.psych.54.101601.145225
- Gabbert, F., Wright, D. B., Memon, A., Skagerberg, E. M., & Jamieson, K. (2012). Memory conformity between eyewitnesses. *Court Review*, 48, 36–43.
- Hall, L., Johansson, P., & Strandberg, T. (2012). Lifting the veil of morality: Choice blindness and attitude reversals on a self-transforming survey. *PLoS ONE*, 7, e45457. doi:10.1371/journal.pone.0045457

- Hall, L., Johansson, P., Tärning, B., Sikström, S., & Deutgen, T. (2010). Magic at the marketplace: Choice blindness for the taste of jam and the smell of tea. *Cognition*, 117, 54–61. doi:[10.1016/j.cognition.2010.06.010](https://doi.org/10.1016/j.cognition.2010.06.010)
- Johansson, P., Hall, L., Gulz, A., Haake, M., & Watanabe, K. (2007). Choice blindness and trust in the virtual world. *Technical Report of the Institute of Electronics, Information, and Communication Engineers—Human Information Processing (IEICE-HIP)*, 107, 83–86.
- Johansson, P., Hall, L., & Sikström, S. (2008). From change blindness to choice blindness. *Psychologia*, 51, 142–155. doi:[10.2117/psysoc.2008.142](https://doi.org/10.2117/psysoc.2008.142)
- Johansson, P., Hall, L., Sikström, S., & Olsson, A. (2005). Failure to detect mismatches between intention and outcome in a simple decision task. *Science*, 310, 116–119. doi:[10.1126/science.1111709](https://doi.org/10.1126/science.1111709)
- Johansson, P., Hall, L., Sikström, S., Tärning, B., & Lind, A. (2006). How something can be said about telling more than we can know: On choice blindness and introspection. *Consciousness and Cognition*, 15, 673–692. doi:[10.1016/j.concog.2006.09.004](https://doi.org/10.1016/j.concog.2006.09.004)
- Johansson, P., Hall, L., Tärning, B., Sikström, S., & Chater, N. (2014). Choice blindness and preference change: You will like this article better if you (believe you) chose to read it! *Journal of Behavioral Decision Making*, 27, 281–289. doi:[10.1002/bdm.1807](https://doi.org/10.1002/bdm.1807)
- Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning and Memory*, 12, 361–366. doi:[10.1101/lm.94705](https://doi.org/10.1101/lm.94705)
- 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, 19–31.
- McLaughlin, O., & Somerville, J. (2013). Choice blindness in financial decision making. *Judgment and Decision Making*, 8, 561–572.
- Meade, M. L., & Roediger, H. L., III. (2002). Explorations in the social contagion of memory. *Memory & Cognition*, 30, 995–1009. doi:[10.3758/BF03194318](https://doi.org/10.3758/BF03194318)
- Merckelbach, H., Jelicic, M., & Pieters, M. (2011). Misinformation increases symptom reporting: A test–retest experiment. *JRSM Short Reports*, 2, 1–6.
- Okado, Y., & Stark, C. L. (2005). Neural activity during encoding predicts false memories by misinformation. *Learning and Memory*, 12, 3–11. doi:[10.1101/lm.87605](https://doi.org/10.1101/lm.87605)
- Pärnamets, P., Hall, L., & Johansson, P. (2015). Memory distortions resulting from a choice blindness task. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th Annual Meeting of the Cognitive Science Society* (pp. 1823–1828). Austin, TX: Cognitive Science Society.
- Roediger, H. L., III, Jacoby, D., & McDermott, K. B. (1996). Misinformation effects in recall: Creating false memories through repeated retrieval. *Journal of Memory and Language*, 35, 300–318.
- Sagana, A., Sauerland, M., & Merckelbach, H. (2014). “This is the person you selected”: Eyewitnesses’ blindness for their own facial recognition decisions. *Applied Cognitive Psychology*, 28, 753–764. doi:[10.1002/acp.3062](https://doi.org/10.1002/acp.3062)
- Sauerland, M., Schell, J. M., Collaris, J., Reimer, N. K., Schneider, M., & Merckelbach, H. (2013). “Yes, I have sometimes stolen bikes”: Blindness for norm-violating behaviors and implications for suspect interrogations. *Behavioral Sciences & the Law*, 31, 239–255. doi:[10.1002/bsl.2063](https://doi.org/10.1002/bsl.2063)
- Tousignant, J. P., Hall, D., & Loftus, E. F. (1986). Discrepancy detection and vulnerability to misleading postevent information. *Memory & Cognition*, 14, 329–338. doi:[10.3758/BF03202511](https://doi.org/10.3758/BF03202511)