Perception & Psychophysics
2004, 66 (8), 1268-1281

Nothing compares 2 views: Change blindness can occur despite preserved access to the changed information

STEPHEN R. MITROFF
Yale University, New Haven, Connecticut

DANIEL J. SIMONS
University of Illinois at Urbana-Champaign, Urbana, Illinois

and

DANIEL T. LEVIN
Vanderbilt University, Nashville, Tennessee

Change blindness, the failure to detect visual changes that occur during a disruption, has increasingly been used to infer the nature of internal representations. If every change were detected, detailed representations of the world would have to be stored and accessible. However, because many changes are not detected, visual representations might not be complete, and access to them might be limited. Using change detection to infer the completeness of visual representations requires an understanding of the reasons for change blindness. This article provides empirical support for one such reason: change blindness resulting from the failure to compare retained representations of both the pre- and postchange information. Even when unaware of changes, observers still retained information about both the pre- and postchange objects on the same trial.

Usually, one's phenomenological experience is of a rich and stable visual world, but where does this sense come from? This richness might arise from a detailed internal representation of the world that can be compared or combined with subsequent views (e.g., Marr, 1982; McConkie & Rayner, 1976; Poggio, Torre, & Koch, 1985; Trehub, 1991). Alternatively, it might arise from the absence of such a representation combined with instantaneous access to the outside world when needed (e.g., Gibson, 1966; Grimes, 1996; Hayhoe, 2000; O'Regan, 1992; O'Regan & Noë, 2002; Rensink, 2000; Stroud, 1955). Over the past decade, change detection tasks have been used extensively to infer the quality and detail of internal visual representations. The primary finding from a variety of experiments is that substantial visual changes can go undetected—a phenomenon known as *change blindness* (for recent reviews, see Hollingworth & Henderson, 2002; Rensink, 2002; Simons, 2000).

In recent years, researchers have explored several causes of change blindness (see Simons, 2000, for a discussion). One explanation suggests that change blindness results from a failure to encode or represent the prechange information (Noë, Pessoa, & Thompson, 2000; O'Regan & Noë, 2002). Unless observers encoded and retained a representation of the initial display, they could not detect a change. Another possibility is that change blindness occurs when adequately encoded information about the initial display is disrupted, overwritten, or forgotten (Beck & Levin, 2003; Becker, Pashler, & Antis, 2000; Brawn, Snowden, & Wolfe, 1999; Irwin, 1992; Levin, Simons, Angelone, & Chabris, 2002; Pashler, 1988; Phillips, 1974; Rensink, O'Regan, & Clark, 1997; Silverman & Mack, 2001; Tatler, 2001; Wolfe, 1999). A third alternative is that change blindness results not from the absence or inadequacy of a representation, but from a failed comparison of a prechange representation to the postchange information (Angelone, Levin, & Simons, 2003; Hollingworth & Henderson, 2002; Ryan & Cohen, 2004; Scott-Brown, Baker, & Orbach, 2000; Shore & Klein, 2000; Simons, 2000; Simons, Chabris, Schnur, & Levin, 2002).

Determining which of these mechanisms contribute to change blindness is central to understanding the completeness of internal representations. If change blindness results entirely from the absence of representations, representations of the visual world might well be sparse,

S.R.M. was supported by a National Science Foundation graduate fellowship and National Institutes of Health Grant F32-MH66553, and D.J.S. was supported by NSF Grant BCS-9905578 and by a fellowship from the Alfred P. Sloan Foundation. We thank Cory Ip and Gisli Palsson for their help with data collection, and George Alvarez, Steven Franconeri, Todd Horowitz, Brian Scholl, Ron Rensink, Daniel Smilek, Jeff Wetherhold, Pamela Yee, and two anonymous reviewers for their helpful comments. We also thank Andrew Hollingworth for suggesting the analysis based on the assumption that observers can retain just one representation and Mike Tarr for providing the image set (http://www.cog.brown.edu/~tarr). Correspondence concerning this article should be addressed to S. R. Mitroff, Department of Psychology, Yale University, P.O. Box 208205, New Haven, CT 06520 (e-mail: stephen.mitroff@yale.edu).

with little or no information retained internally. This view is consistent with the idea that the world itself might act as an "outside memory," with no need to store visual details internally (O'Regan, 1992; O'Regan & Noë, 2002). Alternatively, if change blindness results solely from overwriting of the prechange representation, representations of the visual world might be detailed, but fleeting and fragile. Lastly, if change blindness results from failed comparisons, then representations of the visual world might be detailed and relatively stable even if they are not always readily accessible or cannot be aligned between pre- and postchange views.

The evidence available to date strongly suggests that change blindness, at least in part, can arise from both a failure to form representations (e.g., Noë et al., 2000; O'Regan & Noë, 2002) and a failure to retain adequately formed representations (e.g., Beck & Levin, 2003; Brawn et al., 1999; Rensink et al., 1997). However, strong, direct evidence for failed comparisons of preserved representations is lacking. Here, we present evidence that change blindness can occur even when observers have continued access to information about both the pre- and postchange objects on the same trial. Our evidence suggests that change blindness does not depend entirely on the absence of internal representations—sufficient representations of both the pre- and postchange objects can be formed and retained even in the face of change blindness, suggesting a critical role for a comparison process in successful change detection.

Previous Evidence for Preserved Representations

Several existing lines of research suggest that comparison failures contribute to change blindness (Angelone et al., 2003; Beck & Levin, 2003; Hollingworth, 2003; Hollingworth & Henderson, 2002; Hollingworth, Schrock, & Henderson, 2001; Hollingworth, Williams, & Henderson, 2001; Silverman & Mack, 2001; Simons et al., 2002), but none test this possibility directly; that is, none provide evidence for preserved representations of both the pre- and the postchange objects when the change is not detected. In this section, we briefly discuss the evidence to date for comparison failures and preserved representations.

The first source of evidence for comparison failures comes from a recent study using a cuing technique in conjunction with a change detection task (Hollingworth, 2003). On 50% of the trials, after a change occurred, a cue directed attention to the change location. If change blindness results from a failed comparison of an adequate prechange representation with the postchange display, change detection should be better when the postchange item was cued than when it was not; the cue should allow observers to limit their comparison processes to the target item. In agreement with this prediction, change detection was significantly better on cue trials.

This finding suggests that prechange representations are potentially available for comparison, but the approach has several drawbacks as a way of inferring comparison failures. First, the measures of preserved representation rely on successful change detection; any inferences about preserved representations are based solely on trials with successful detection and not on those in which changes are unreported. Second, due to the dependence on successful detection trials, the design does not completely decouple preserved representations and the comparison process; preserved representations must be inferred through successful comparisons. Finally, the cuing approach detects preserved representations of the prechange display only because the postchange display remained visible at test. Therefore, there is no way of knowing whether change blindness has occurred because of a representational failure in which a representation of one object (in this case, the prechange object) took precedence at the expense of the other.

A second line of evidence for preserved representations despite change blindness comes from a collection of video-based and real-world change blindness experiments (Angelone et al., 2003; Simons et al., 2002). In one study, observers watched a video in which a change occurred across a camera cut (Angelone et al., 2003). Even the observers who failed to notice the change were able to select the prechange object out of a photographic lineup at above-chance levels, suggesting that they had a preserved representation of the prechange information. In addition, observers who failed to detect the change were just as accurate in recognizing the prechange object as were those who detected the change. In another study, an experimenter carrying a red-and-white striped basketball approached a pedestrian (the subject) and asked for directions to the campus gymnasium (Simons et al., 2002). While the pedestrian provided directions, a group of people walked by and surreptitiously took the ball. Although most subjects did not report noticing the change, when they were subsequently asked directed questions about what the experimenter had been carrying, most recalled the basketball and could even describe its unusual color pattern. Again, observers retained a representation of the prechange information even when they did not notice the change.

These findings unquestionably provide evidence for preserved representation of the prechange information, but their implications for comparison failures are less clear. First, the single-trial design of these incidental change detection tasks does not allow for a systematic analysis of what information is and is not available when the change is and is not detected. Second, neither of these studies examined whether or not the observers had a preserved representation of the postchange actor or object (the postchange display was still visible at test). In fact, a third study exploring this issue found no evidence for a preserved representation of either the pre- or the postchange features when observers did not notice the change (Levin et al., 2002). Given the single-trial design and the variability inherent in real-world studies, the cause of this null effect is unclear.

A third line of research exploring preserved representations used simple arrays of objects and examined both pre- and postchange representations (Beck & Levin, 2003). When one object in an array was replaced by a new object, observers could successfully select the postchange object but not the prechange object from a lineup. Interestingly, even when there was no postchange object (the prechange item just disappeared), if the second display underwent a slight shift to the left or right, observers were still poor at selecting the prechange object. Only when there was no postchange object and no shift of the second display did observers successfully recognize the prechange object (Beck & Levin, 2003). Thus, this study provides little evidence for preserved representations, suggesting instead that change blindness might result primarily from overwriting. Despite its more systematic approach, this study has a drawback that complicates the interpretation. The study was designed to obtain an uncontaminated measure of the representation of the initial object. Consequently, observers were not asked whether or not they detected the change. Thus, the critical question of whether or not observers were able to form a representation of the pre- and postchange objects when they failed to detect the change was not addressed by this experiment.

Finally, a fourth source of evidence for preserved representations comes from an experiment that directly tested change detection performance and also measured both pre- and postchange representations (Silverman & Mack, 2001). Subjects performed a modified Sperling task in which one row of letters in a 3×3 array changed. When observers reported that they had not seen a change, they still showed priming for the pre- and postchange rows in a subsequent task. Although this study suggests that change blindness might result from a comparison failure, such an inference is not entirely warranted because the pre- and postchange representations were tested in a between-groups design; one group of subjects showed priming for the prechange row, and a different group showed priming for the postchange row. Consequently, the study does not test whether observers had preserved representations of both the pre- and postchange objects when they failed to detect the change. The change might have been missed because observers represented only the pre- or the postchange object on a given trial.

Summary

Taken together, these four lines of research suggest that a representation of the prechange object is available even when a change is not detected, that this representation can be retained across a number of subsequent fixations, and that it can be compared with a postchange display for successful change detection. However, none of these experiments provide evidence that representations of the pre- and postchange displays are retained when changes are and, more important, *are not* detected. This question is central to understanding whether representation failures underlie change blindness. If we can obtain evidence for sufficiently preserved representations of both the pre- and postchange objects in a single instance of change blindness, the possibility of a comparison failure is considerably strengthened. Note that such a comparison failure explanation does not require that the representations be perfect or complete, only that they need to contain sufficient information to potentially support change detection.

The four experiments reported in this article are based on those of Beck and Levin (2003) in which observers viewed paired arrays of objects. The goal of our study was to explore what information was preserved on a given trial when observers failed to notice a change. On each trial, an array appeared, disappeared, and then reappeared with one of the objects potentially replaced by a new one. Observers reported whether or not they detected a change and answered a series of two-alternative forced-choice (2AFC) questions. For the 2AFC questions, observers were shown two objects and were asked to decide which one had been in the displays on the trial. These questions were designed so that observers were forced to choose between the prechange object and a novel object on one question, and between the postchange object and a novel object on another. Accuracy on these 2AFC questions served as our operational measure of *sufficiently* retained representations; that is, if observers retained enough information, they should select the object that had been present.

EXPERIMENT 1

This experiment explored whether or not observers could simultaneously retain sufficient representations of both the pre- and postchange information when they failed to detect the change. Since both pre- and postchange representations were examined on each trial, we could determine whether or not observers are better than chance at correctly selecting both the pre- and postchange objects in the 2AFC questions; is the proportion of trials where observers were accurate on both the pre- and postchange questions greater than would be expected by chance? Answering this question requires a definition of what constitutes chance performance. As discussed below, we compared performance with two levels of chance, a lower boundary level that estimated chance performance if the observers simply guessed on each trial (retained no representations) and an upper boundary level that estimated the absolute maximum performance possible if observers were unable to simultaneously retain sufficient representations of both the pre- and postchange objects. If change blindness results from a comparison failure and observers retained sufficient representations of both the pre- and postchange objects even when they failed to detect a change, the probability of answering both the pre- and postchange 2AFC questions correctly when unaware of the change should exceed both of these calculations of chance performance.

Method

Subjects. Seventeen individuals were paid for a single 1-h session. They were encouraged to be liberal when reporting change detection, in order to minimize the likelihood that unaware trials (trials for which no change was reported) would include mislabeled aware trials (trials for which a change was reported). However, too liberal a criterion might signify that the observers were not performing the task appropriately, so data were eliminated for observers who had a false alarm rate more than two standard deviations above the mean. Data from 1 observer who surpassed this criterion (reporting changes on 63% of the no-change trials) were eliminated, and all further analyses were based upon data from the remaining 16 observers.

Apparatus. Stimuli were presented on a Macintosh iMac computer using custom software written with Vision Shell C libraries (http://www.visionshell.com). The observers sat comfortably without head restraint approximately 50 cm from the monitor.

Materials. For each trial, a display consisted of 4, 6, or 8 line-drawn objects randomly chosen without replacement from a set of 260 line-drawn objects (Snodgrass & Vanderwart, 1980). The objects were randomly located in an imaginary 4 × 4 grid with a uniform white background (see Figure 1). The entire grid was framed by a 1-mm-thick black line and was approximately 23º × 23º. The objects were digitized in 256-level grayscale and were presented at 72 dpi. Although the objects varied in size, they were a maximum of 4.5º in either dimension. The objects for the 2AFC questions were presented side by side in an 12.6º × 6.3º rectangle, with a 1-mm-thick black surround (see Figure 1).

Procedure. On each trial, an array of 4, 6, or 8 line-drawn objects appeared for 1,000 msec, was replaced by a blank white display of the same luminance as the background for 350 msec, and then reappeared for an additional 1,000 msec. When the array reappeared, the entire grid was shifted 0.8º randomly to the left or right. This shift created a change signal at each object location to reduce the effect of any transient signal from the changed object that might survive the blank interval (Beck & Levin, 2003). On 90% of the trials (change trials), one of the objects was replaced by a new object selected at random from the remaining objects of the entire set. The no-change trials (10%) were included to provide a measure of false alarms, and the three set sizes were included to produce a range of change detection performance. After the second display had disappeared, the observers responded to four discrete questions as quickly as possible (three 2AFC questions and one change detection question).

2AFC questions. Three sequential 2AFC questions were used to assess the observers' memory for the objects that had been present on the trial. For each question, the observers were shown two objects, one that had been present in the displays of that trial and one that had not. Their task was to select the object that had been present. For the *prechange question*, one object was the prechange object, and the other was a novel object that had not been present on the trial. For the *postchange question*, one object was the postchange object and the other was a novel object. For the *unchanged object question*, one object was an unchanged object from the displays, and the other was a novel object. The unchanged object used for this question was randomly chosen from the set of unchanged objects on the trial. The novel objects for all three questions were randomly selected from the remaining objects from the entire set. The three questions were presented in a random order for each trial. On no-change trials, the observers responded to three 2AFC questions that referred to one object that had been in the displays from the trial and one novel object. None of the objects were repeated

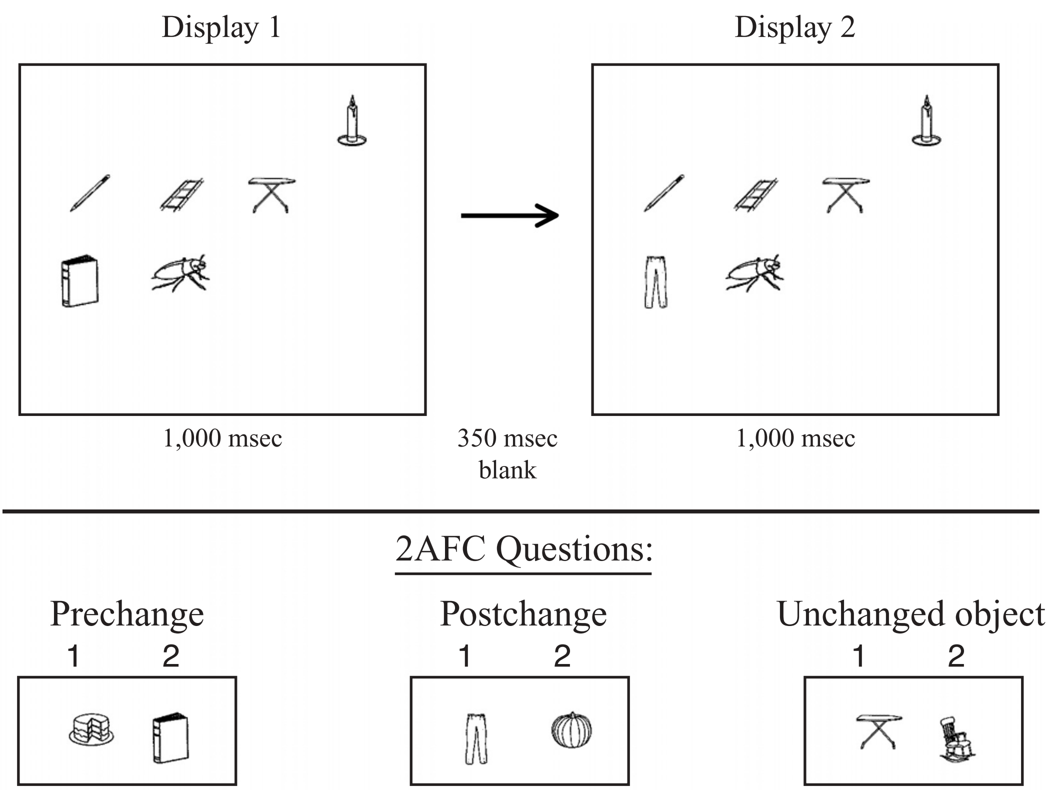

Figure 1. In each experiment, line-drawn objects were randomly positioned in a 4 × 4 imaginary grid (Display 1). On change trials, when the array reappeared after a 350-msec blank display, one object had changed (Display 2). The prechange, postchange, and unchanged object 2AFC questions were presented in a random order. Observers reported which of the two objects from each question was present on the trial.

across the three no-change trial questions. For each question, the novel object was randomly placed on the left or right, and observers pressed "1" or "2" to signify that the left or right object (respectively) had been in the displays on the trial.

The observers were told of the composition and random presentation order of the three 2AFC questions, and it was emphasized that their only concern was to select the object that had been present. They did not need to know which object was, for example, the prechange object; they only needed to report which object from each question had been in the displays of the trial. The observers were instructed to guess when they did not know which object had been present.

Change detection question. After the 2AFC questions, the observers reported whether or not they had detected a change on the trial. They were given the following instructions, designed to induce a liberal criterion for a detection response: "For a response of *yes*, you need not have complete knowledge of the change (i.e., the identity of the pre- and postchange objects) but you should have knowledge about its location." They were also told to base their change detection response on their original impression from the displays and not to try to infer the presence of a change from the content of the 2AFC questions. Observers were informed of the no-change trials and instructed to try to avoid making a large number of false alarms; we wanted the observers to adopt a liberal criterion and to respond *yes* if they had any knowledge about the occurrence of a change, but we did not want them to respond haphazardly (observers typically adopt conservative criteria in change detection tasks, reporting changes only when relatively certain about the presence of a change). The observers pressed the "1" and "2" keys for a *yes* and *no* response, respectively.

Following verbal instructions, the observers completed 15 practice trials and 100 test trials of each set size, with 10 of each 100 being no-change trials (300 total test trials). All of the trials were presented in a unique random order for each observer. Any questions were answered after the practice trials and before the test trials. The observers were allowed to take breaks between any two trials when needed, and after every 50 trials, a message box appeared, encouraging them to take a short break.

Results and Discussion

The change detection question and the 2AFC questions together allow an assessment of the presence of pre- and postchange object representations when observers did and did not report noticing a change.

Change detection. On the basis of a signal detection analysis, the observers were sensitive to the presence of changes [mean $A' = .825$, $SD = .074$; $t(15) = 17.58$, $p < .001$][1] and were biased to report no change [mean $B'' = .273$, $SD = .359$; $t(15) = 3.05$, $p = .008$]. Collapsed across set size, the observers reported 65.88% (SD = 15.38%) of the changes and falsely reported changes on 18.54% (SD = 14.35%) of the no-change trials. As set size increased, reports of change detection decreased [for set size 4, M = 85.69%, SD = 14.55%; set size 6, M = 64.10%, SD = 17.66%; set size 8, M = 47.85%, SD = 17.96%; $F(2,30) = 85.53$, $p < .001$]. False alarm rates, however, remained stable [for set size 4, M = 18.75%, SD = 16.68%; set size 6, M = 15.63%, SD = 17.88%; set size 8, M = 21.25%, SD = 18.57%; $F(2,30) = 0.78$, $p = .467$].

2AFC questions. Accuracy on the 2AFC questions provides a measure of memory for the prechange, postchange, and unchanged objects on each trial (response time data are provided in Appendix A). For example, if, on average, observers correctly select the prechange object over the novel object on the prechange 2AFC question, we will infer that they have a sufficiently preserved representation of the prechange object. As shown in Figure 2, observers were significantly better than chance (50%) at selecting the prechange, postchange, and unchanged objects in the 2AFC questions both when aware and unaware of the change (all t values > 3.87, all ps $< .002$).[2]

Were both the pre- and postchange objects retained on the same trial more often than expected by chance? On average, observers correctly selected both the pre- and postchange objects on the same trial 72.24% (SD = 10.18%) of the time when aware of the change and 42.98% (SD = 9.59%) when unaware. To determine whether these levels of memory performance might indicate the existence of comparison failures, we compared performance on unaware trials with two different levels of chance (see Table 1).

Lower boundary of chance. If we make no a priori assumptions about the possibility of retaining pre- and postchange representations, the probability of happening to be correct on both the pre- and postchange 2AFC questions is the product of the individual probabilities. If we assume that subjects simply guessed on each question, they should be correct on each question 50% of the time. Thus, performance based on guessing alone would predict correct responses on both questions on 25% of the trials. Observed performance easily exceeded this estimate when observers were unaware of the change [$t(15) = 7.50$, $p < .001$]. However, this null hypothesis does not provide a sufficiently rigorous test of the claim that change blindness must result, in part, from a comparison failure. Although a correct response to both questions could arise because of correct guesses on both questions, it could also arise from accurate memory on one question and guessing on the other. If so, performance might exceed 25% even if observers failed to represent both the pre- and postchange objects on the same trial.

Upper boundary of chance. The strongest test of the comparison failure hypothesis depends on assuming that observers can represent at most either the pre- or postchange object, but not both. That is, if we assume that both representations cannot be retained, what is the absolute maximum percentage of trials where the observers will be able to guess correctly on one or both of the pre- and postchange questions? If performance exceeds this estimate, change blindness would have occurred on some trials even when observers had sufficient representations of the pre- and postchange objects. Note that this estimate of chance is *extremely* conservative, maximizing the effectiveness of memory for either the pre- or postchange object, while excluding the possibility of representing both. Consequently, we would not expect performance to exceed this estimate by much. Provided that performance is reliably greater, the finding would provide strong evidence for a comparison failure.

To calculate this upper boundary of chance, for each observer, we estimated the percentage of trials where they (1) retained a prechange representation and guessed

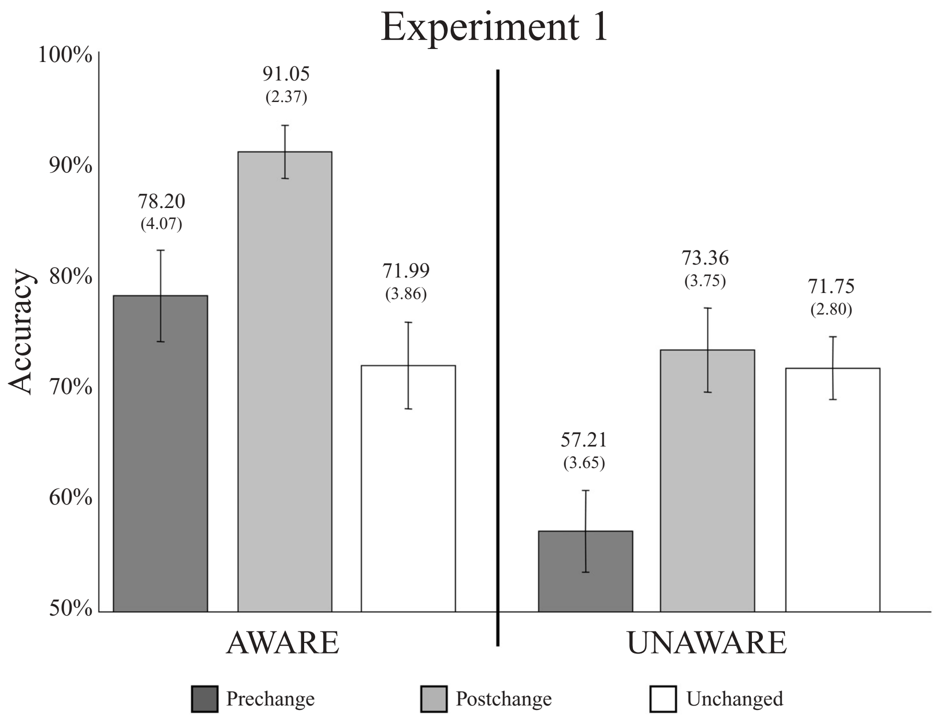

Figure 2. Accuracy for the prechange, postchange, and unchanged object 2AFC questions of Experiment 1. Means (and standard deviations, in parentheses) are presented above the corresponding bars. Error bars represent 95% confidence intervals, and all means are significantly above chance (50%). Aware trials are presented on the left and unaware trials on the right.

correctly on the postchange question, (2) retained a postchange representation and guessed correctly on the prechange question, and (3) retained neither a pre- nor a postchange representation and guessed correctly on both questions. The sum of these three components provides the absolute maximum percentage of trials where observers will respond correctly to both the pre- and postchange questions without having retained representations of both. For each observer, we must first isolate the probability of retaining a representation of the prechange object alone and of the postchange object alone. That is, we need to know what percentage of the correctly answered questions was based on a retained representation and what percentage resulted from happening to guess correctly. Given that observers are equally likely to be correct or incorrect when guessing on a 2AFC question, we can estimate the percentage of correct trials resulting from accurate guesses from the percentage of trials where the observers were incorrect; the percentage of accurate guesses should equal the percentage of inaccurate guesses. For example, take a hypothetical observer who was accurate on the prechange question 58% of the time

Table 1
Percentage of Trials in Which Observers Correctly Answered Both the Pre- and Postchange 2AFC Questions When Reporting No Awareness of the Change for Experiments 1, 2, and 3

Unaware Trials	Actual Value	Chance, Lower Boundary	Chance, Upper Boundary
Experiment 1	42.98%	25.00% $t(15) = 7.50, p < .001$	40.63% $t(15) = 3.43, p = .004$
Experiment 2	48.45%	25.00% $t(11) = 12.48, p < .001$	44.48% $t(11) = 5.57, p = .001$
Experiment 3	42.70%	25.00% $t(11) = 5.82, p < .001$	40.32% $t(11) = 2.55, p = .027$

Note—The values are compared with two levels of chance: a lower boundary, which makes no assumptions about the observers' ability to retain a representation, and an upper boundary, which assumes that observers can retain, at most, only a single representation on a given trial (calculated from Formula 1).

and the postchange question 70% of the time when unaware of the change. For the prechange question, this observer guessed *incorrectly* 42% of the time and presumably guessed *correctly* 42% of the time. If we eliminate the correct guesses, the remaining 16% of the trials reflect correct answers based on a retained representation. Likewise for the postchange question, this observer guessed incorrectly (and thus also correctly) 30% of the time, so 40% of his or her correct answers reflect a retained representation.

The critical assumption of this chance estimate is that observers can only retain a representation of one object or the other. Thus, when they represent the prechange object, they have to guess on the postchange question. When guessing, they will be correct 50% of the time. Likewise, when they retain the postchange object, they must guess on the prechange question, and they will be correct on 50% of their guesses. For our hypothetical observer, the probability of retaining a representation of the prechange object and guessing correctly on the postchange object question is 16% * 50% = 8%. For the postchange question, this probability is 40% * 50% = 20%. Thus, this observer could retain a representation of either the pre- or postchange object on a given trial and guess on the other question, leading to correct answers on both questions, on 28% of the trials. Finally, even when neither the pre- nor postchange object was retained, observers would guess on both questions, and they would be correct on both 25% of the time. Our hypothetical observer retained the prechange object on 16% of the trials and the postchange on 40%, which leaves 44% of the trials for which neither was retained. The probability of answering both questions correctly for these trials is simply 25% of this 44% of trials, so on 11% of the trials, the observer would guess correctly on both the pre- and postchange questions. Adding these trials in which both questions were guessed correctly to the trials in which one object was represented and the other was guessed correctly gives the maximum proportion of trials for which both questions could be answered correctly without representing both objects.

Specified formulaically, if P is the proportion of trials for which both questions are answered correctly without representing both the pre- and postchange objects, and R is the percentage of trials for which an object was retained,

$$P = 50\% * (R_{pre} + R_{post}) + .25\% * [100\% - (R_{pre} + R_{post})]. \quad (1)$$

P was calculated separately for all observers on the basis of their separate 2AFC accuracy on the pre- and postchange questions. We then compared the actual proportion of trials for which that observer answered both the pre- and postchange questions correctly on the same trial with P. If observers can simultaneously retain representations of both the pre- and postchange objects, they should be correct on both questions more frequently than P. Collapsed across set size, observers performed significantly above P when unaware of the change [mean percentage = 42.98%; mean P = .4063, *SD* = .0733; paired $t(15) = 3.43, p < .004$]. Given that P is an extremely conservative measure of chance performance and most likely overestimates the proportion of correct responses to both questions that resulted from representing only one object, the finding that performance significantly exceeded P represents a strong test of the hypothesis that observers actually represent both the pre- and postchange objects when they fail to detect a change.

Failed comparisons. The preserved representations for both the pre- and postchange objects on unaware trials show that change blindness occurs, at least in part, from failed comparisons. Specifically, when unaware of the change, observers answered both the pre- and postchange object 2AFC questions at better than chance levels. More important, performance was better than would be expected on the basis of the strongest assumption that at most a single representation could be retained. Thus, when observers have access to the necessary representations, they still can fail to make the necessary comparison. This experiment provides the first conclusive evidence that observers can access both the pre- and postchange information on a given trial but still fail to notice a change.

Although this finding is consistent with the idea that change blindness results from a comparison failure, it is also possible that the representations involved in recognition tasks (i.e., the 2AFC questions) are not the same as those involved in detection tasks (i.e., the change detection question); different representations might underlie the two tasks. In other words, the recognition task might not be an exhaustive measure of all possible representations. If it is not, change blindness could still result from a representation failure, but the form of representation underlying the 2AFC and change detection tasks would have to be separate. If different types of representation underlie performance on these two tasks, it is unlikely that performance on one task will correlate with performance on the other. Alternatively, if one representation underlies performance in both tasks, change detection performance should be related to recognition performance. In the present experiments, as well as others (i.e., Levin et al., 2002), observers performed better on the recognition tasks when they detected the changes. For each observer, we calculated a correlation between change detection and accuracy on (1) the prechange question, (2) the postchange question, and (3) both. Change detection performance was positively correlated with accuracy on the prechange question [mean $r = .219$; $t(15) = 8.97$, $p < .001$],[3] the postchange question [mean $r = .236$; $t(15) = 11.65$, $p < .001$], and both the pre- and postchange questions [mean $r = .291$; $t(15) = 10.39$, $p < .001$]. This finding strengthens the claim that change blindness can result, in part, from a comparison failure.[4]

Finally, even if the representations underlying the recognition and detection tasks are related, better performance on one than the other might just reflect differences in the sensitivity of the measures. That is, 2AFC might be a more sensitive measure than change detection. Although the present experiments are not suited to address this issue, future studies could adopt more sensitive measures of change detection (e.g., eye movement patterns) to offset this possibility. The present experiments show that representations of both the pre- and postchange objects may coexist even when the most commonly used measure of change detection (explicit reports) reveals change blindness.

Summary

Experiment 1 provides evidence that when observers had access to representations of both the pre- and postchange objects, they still missed some changes, suggesting that change blindness can occur from a comparison failure. However, we must consider the possible artifactual explanation that observers knew that the set of three 2AFC questions involved the object that changed and another object from the display, so that when viewing the arrays of objects, they might have attempted to remember as many of the objects as possible at the expense of change detection performance. That is, observers might have shirked their change detection responsibilities to encode more information about individual objects. If observers used such a strategy, the simultaneous retention of both the pre- and postchange objects on unaware trials might be a consequence of the task demands. The finding that recognition performance was improved when changes were detected somewhat mitigates this concern, but Experiment 2 addresses it directly by placing a greater emphasis on the change detection task.

EXPERIMENT 2

In Experiment 2, we focused attention on the change detection task by eliminating the unchanged object 2AFC question and informing observers that the 2AFC questions would always contain objects involved in the change. Given this modification, the task demands should lead observers to focus all of their attention on trying to find the change. If the above-chance performance on both the pre- and postchange 2AFC questions from Experiment 1 resulted from a strategy of neglecting the change detection task in favor of trying to encode all of the objects in each display, it should be lessened or eliminated by this procedural modification, and change detection performance should increase. Alternatively, if the probability of responding correctly on both pre- and postchange object 2AFC questions of Experiment 1 was unaffected by this task strategy, it should remain above chance.

Method

Subjects. Thirteen individuals were paid for a single 1-h testing session. One observer's data were eliminated because she misunderstood the instructions on the 2AFC questions and selected the item that was not present in the displays of the trial. No data were eliminated by applying the false alarm criteria from Experiment 1, and all analyses were based on data from 12 observers.

Procedure and Materials. Experiment 2 was identical to Experiment 1, except for the following three differences: (1) All 300 trials were six-object displays, (2) there were 75% change trials and 25% no-change trials, and (3) the unchanged object 2AFC question was eliminated. A single set size was used to allow for a more powerful assessment of performance on change trials, and a greater number of no-change trials was used to obtain more precise measures of sensitivity and bias.

Results and Discussion

Change detection. As in Experiment 1, observers were sensitive to the changes [mean $A' = .851$, $SD = .041$; $t(11) = 29.67$, $p < .001$] and biased to report no change [mean $B'' = .332$, $SD = .341$; $t(11) = 3.37$, $p = .006$]. There were no significant differences between Experiment 1 and the present experiment for either sensitivity [$t(26) = 1.18$, $p = .248$] or bias [$t(26) = 0.44$, $p = .663$]. On average, observers detected 66.00% ($SD = 12.45\%$) of the changes and reported a change on 14.11% ($SD = 12.45\%$) of the no-change trials.

2AFC questions. As can be seen in Figure 3A, observers responded correctly on the pre- and postchange object questions more than 50% of the time, both when aware and unaware of the change (all t values > 6.17, all ps $< .001$). Observers were accurate on both the pre- and postchange 2AFC questions simultaneously on 72.28% ($SD = 7.33\%$) of the trials when aware of the change and on 48.45% ($SD = 6.51\%$) of the trials when unaware. As for Experiment 1, the simultaneous performance on the unaware trials was compared with both a lower and an upper boundary of chance (see Table 1). With no a priori assumptions about what can or cannot be represented, chance accuracy on both questions for a given trial is 25%. Observers performed significantly above this level when unaware of the change [$t(11) = 12.48$, $p < .001$]. When we conservatively assumed that at most a single representation can be retained on a given trial (i.e., of either the pre- or the postchange object, but never both) and used Formula 1 from Experiment 1, P for unaware trials was 44.48% ($SD = 4.90\%$). Observers were significantly more accurate on both the pre- and postchange questions than P [$t(11) = 5.57$, $p < .001$].

As in Experiment 1, change detection performance was positively correlated with accuracy on the prechange question [mean $r = .135$; $t(11) = 4.44$, $p = .001$], the postchange question [mean $r = .277$; $t(11) = 10.28$, $p < .001$], and both the pre- and postchange questions [mean $r = .235$; $t(11) = 8.19$, $p < .001$]. Thus the results of the present experiment replicate those of Experiment 1; even when observers fail to detect the change, they can retain representations of both the pre- and postchange objects. For an analysis of the order of presentation of the pre- and postchange questions, see Appendix B.

Summary

Experiment 2 replicated and strengthened the finding that observers can have access to information about the pre- and postchange objects when the change goes undetected.

EXPERIMENT 3

Experiments 1 and 2 showed that observers can have access to sufficiently preserved representations of both the pre- and postchange objects while failing to detect

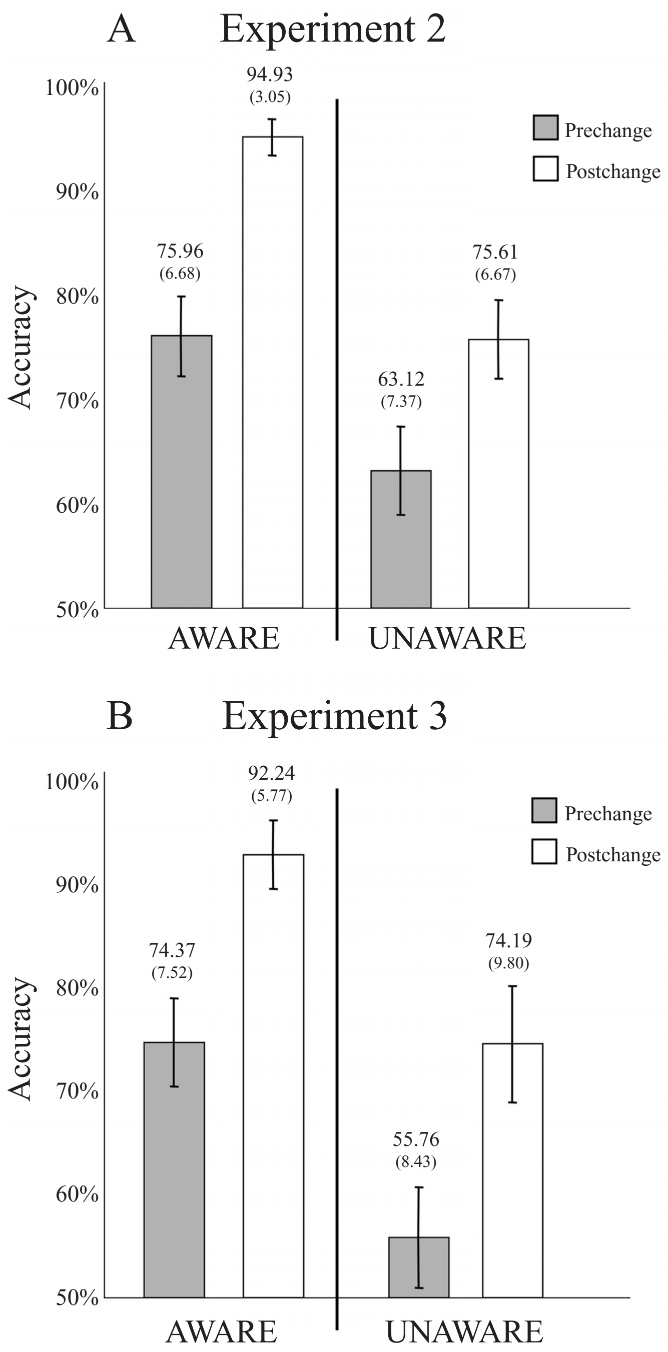

Figure 3. Accuracy for the pre- and postchange 2AFC questions of Experiments 2 and 3 are presented in panels A and B, respectively. Means (and standard deviations, in parentheses) are presented above the corresponding bars. Error bars represent 95% confidence intervals, and all means are significantly above chance (50%). Aware trials are presented on the left side of each panel and unaware trials on the right.

the change, but they also suggest that the prechange representation might not be as stable as the postchange representation. Observers consistently performed better on the postchange object 2AFC question than on the prechange question. This finding is consistent with previous claims that prechange representations can be disrupted or overwritten either by postchange information or by other aspects of subsequent displays (Beck & Levin, 2003; Brawn et al., 1999). Here we explore the durability of the prechange representation by making a single change to the experimental design. The two tasks of Experiments 1 and 2 (the 2AFC questions and the change detection question) were optimally ordered for determining which representations were available; immediately after the change presentation, observers answered the 2AFC questions. In Experiment 3, observers answered the change detection question first. If the prechange representation is robust, performing the change detection task prior to answering the 2AFC questions should have no adverse effect on accuracy. Alternatively, if the representation is somewhat fragile, accuracy might suffer.

Method

Subjects. Sixteen individuals were paid for a single 1-h testing session. Data were eliminated from 1 observer who had a false alarm rate (51%) more than two *SD*s above the mean. Three observers were forced to evacuate the building due to a fire and failed to complete the experiment. All analyses are based on data from the remaining 12 observers.

Procedure and Materials. Experiment 3 was identical to Experiment 2, except that the change detection question was asked prior to the 2AFC questions.

Results and Discussion

Change detection. As in Experiments 1 and 2, observers were sensitive to the changes [mean $A' = .863$, $SD = .079$; $t(11) = 15.87, p < .001$] and biased to report no change [mean $B'' = .237$, $SD = .276$; $t(11) = 2.98$, $p = .013$]. There was no significant difference between the three experiments for either sensitivity [$F(2,37) = 1.141, p = .330$] or bias [$F(2,37) = 0.25, p = .778$]. Observers reported detecting a change on 71.11% (SD = 15.02%) of the change trials and on 14.00% (SD = 8.19%) of the no-change trials.

2AFC questions. Accuracy values are reported in Figure 3B, and as in Experiments 1 and 2, observers were reliably above 50% on the pre- and postchange questions when aware and unaware of the change (all t values > 2.36, all ps < .038). Observers were accurate on both questions simultaneously on 69.47% (SD = 9.89%) of the trials when aware and 42.70% (SD = 10.53%) when unaware. The proportion of trials for which observers were accurate on both the pre- and postchange questions when unaware of the change was compared with the two levels of chance (see Table 1). As in Experiments 1 and 2, the unaware value was significantly above both the lower boundary level of chance [$t(11) = 5.82, p < .001$] and the conservative upper boundary level [P = .4032, $SD = .0829$; $t(11) = 2.55, p = .027$]. Change detection performance was again positively correlated with accuracy on the prechange question [mean $r = .170$; $t(11) = 9.67$, $p < .001$], the postchange question [mean $r = .244$; $t(11) = 10.34$, $p < .001$], and both the pre- and postchange questions [mean $r = .240$; $t(11) = 12.34$, $p < .001$].

The specific question for this experiment was whether or not accuracy on the prechange question would be adversely affected by delaying it until after the change detection question. We first examined whether any differences existed between Experiment 2 and Experiment 3. A 2 × 2 × 2 analysis of variance (ANOVA) was calculated with two within-subjects variables: *awareness* (aware vs. unaware trials) and *prechange–postchange* (pre- or postchange question accuracy), and one between-subjects variable (Experiment 2 vs. Experiment 3). Neither the three-way interaction nor the two-way interactions revealed significant differences between the two experiments. However, our a priori question of interest was whether or not performance on the prechange question would differ between the two experiments. For the aware trials, there was no significant difference in accuracy between experiments for the prechange question [$t(22) = 0.55, p = .589$]. However, for the unaware trials, accuracy on the prechange question was significantly reduced in Experiment 3, relative to Experiment 2 [$t(22) = 2.27$, $p = .033$]. When the change detection question was asked prior to the 2AFC questions, accuracy for the unaware prechange object question was adversely affected. Note, however, that this result might reflect a speed–accuracy tradeoff. Response times decreased slightly for the unaware prechange condition relative to Experiment 2 (M = 80.49 msec), but they increased slightly for the other three conditions (M = 54.04 msec). However, none of these response time differences were statistically significant. Response times are provided in Appendix A.

Given the detrimental effect of the order of questioning on the prechange representation when the observer was unaware of the change, it is curious that no order effect was found for the prechange question within the 2AFC questions (see Appendix B). In all three experiments, accuracy on the prechange question was unaffected by the order of the 2AFC questions. If the representation was so fragile that the change detection question disrupted it, why was it unaffected by the postchange question? One possibility in need of further research is that the representation is stable enough to withstand only one type, or nature, of questioning; that is, as long as the questions are of a similar form and do not require a different style of thought or response, the representation remains somewhat stable.

Summary

A comparison of Figures 3A and 3B reveals only one consequence of reversing the question order in the present experiment: Accuracy on the prechange object question was worse when observers were unaware of the change. This effect suggests that when observers fail to detect a change, they have only a fragile representation of the prechange object. In combination with Experiments 1 and 2, the present results strongly suggest that observers can simultaneously retain representations of both pre- and postchange information, even when unaware of the change. However, one additional concern remains; that is, all three experiments involved a dual-task situation, and it is unknown how much the 2AFC questions might have affected change detection performance. Experiment 4 addressed this concern.

EXPERIMENT 4

To address the potential effects of the dual-task nature of Experiments 1, 2, and 3 on change detection performance, the 2AFC questions were removed. If answering the 2AFC questions adversely affected change detection accuracy, change detection should increase and false alarms should decrease in Experiment 4. Alternatively, if observers were performing the change detection task adequately, removing the 2AFC questions should have little effect.

Method

Subjects. Thirteen individuals were paid for a single 45-min testing session. Due to a computer malfunction, 1 observer was unable to complete the experiment. All analyses were based on the data of the remaining 12 observers.

Procedure and Materials. Experiment 4 was identical to Experiment 3, except that no 2AFC questions were asked. Observers merely responded "yes" or "no" to the change detection question.

Results and Discussion

As in the other experiments, the observers were sensitive to the changes [mean $A' = .856$, $SD = .041$; $t(11) = 29.91$, $p < .001$] and biased to report no change [mean $B'' = .519$, $SD = .299$; $t(11) = 6.01$, $p < .001$]. Across the four experiments, there was no significant difference in either sensitivity [$F(3,48) = 1.00$, $p = .405$] or bias [$F(3,48) = 1.85$, $p = .151$]. Observers reported detecting a change on an average of 60.70% ($SD = 12.23\%$) of the change trials and 9.11% ($SD = 7.21\%$) of the no-change trials. The dual-task nature of the earlier experiments did not significantly affect change detection performance.

GENERAL DISCUSSION

These four experiments provide evidence for the preservation of representations, even under conditions of change blindness. Experiment 1 showed that observers, whether they did or did not detect a change, sometimes encoded and retained sufficient representations of both the pre- and postchange objects. In Experiment 2, the preservation of the pre- and postchange representations was replicated when a greater emphasis was placed on the change detection task. Experiment 3 confirmed that the prechange representation might be relatively weaker than the postchange representation when the change goes undetected, and Experiment 4 showed that the dual-task nature of Experiments 1, 2, and 3 had no significant effect on change detection performance.

Change blindness can result, at least in part, from the absence of sufficient representations (e.g., Noë et al., 2000; O'Regan & Noë, 2002). In isolation, such work might suggest that internal representations were not formed. However, change blindness might also result, at least in part, from the failure to retain a representation after forming it (e.g., Beck & Levin, 2003; Brawn et al., 1999; Rensink et al., 1997). Accordingly, change blindness sometimes occurs because representations are fleeting or fragile. The present study is the first to demonstrate directly that change blindness also occurs even when observers can have access to sufficient representations of both the pre- and postchange information, suggesting that change blindness can result, at least in part, from the failure to compare representations.

Experiments 1–3 showed that when observers are unaware of a change, they might still have access to representations of both the pre- and the postchange information. One concern about these findings, however, is that these "unaware" trials might be contaminated by mislabeled "aware" trials. That is, even though we encouraged our observers to adopt a liberal criterion for reporting a change (and they had relatively high false alarm rates), they might have sometimes mislabeled a trial as "unaware" when they actually had detected the change. If a few observers adopted too conservative a criterion, their contaminated "unaware" trials might bring the average results above the calculated chance levels. Three additional findings ameliorate this concern. First, if the above-chance recognition findings for the unaware trials resulted from contamination caused by some observers adopting too conservative a bias, those observers with a more conservative bias should have had a higher proportion of trials where they responded correctly to both the pre- and postchange 2AFC questions when unaware. Yet, across all three experiments, the conservativeness of the observers' bias was unrelated to their simultaneous accuracy on the pre- and postchange questions when unaware of the change ($r = .081$). Second, it is unlikely that a few observers were driving the results, since 34 of the 40 observers had a greater value for their actual simultaneous accuracy on both the pre- and postchange 2AFC questions than on their calculated level of chance (P). Finally, across the experiments, 6 observers provided unequivocal evidence for simultaneous retention of representations of both the pre- and postchange objects when unaware of the change; that is, adding the trials where they correctly retained a representation of the prechange object to the trials where they correctly retained a representation of the postchange object results in a sum greater than 100%. By definition, they must have simultaneously retained sufficient pre- and postchange object representations on some trials even when they reported no change.

Although variations in bias across subjects do not appear to account for the simultaneous representation of the pre- and postchange objects, it is possible that a few mislabeled trials could account for better than expected recognition performance for a given observer. If, on a few trials, observers had some access to the change, but not enough to say they noticed it, those trials could contribute to better than chance recognition performance. That is, observers might still have access to both the pre- and postchange representations only because they reported no change but actually detected the change at some level. Although we cannot eliminate this possibility, our study does show that observers have access to both the pre- and postchange representations even when they subjectively report no awareness of the change. Explicit reports of change detection have served as the basis for most evidence of change blindness. As such, any evidence of preserved representations of both the pre- and postchange representations in the face of failed explicit reports would suggest that those explicit failures might result from failed comparisons. Experiments 1–3 represent the strongest evidence to date that change blindness can occur despite continued access to information about the change.

Taken as a whole, the present experiments suggest that multiple representations of the outside world are stored internally and that these representations can be disrupted by later events. However, a few open issues remain. Namely, the results do not speak directly to the *nature* or *completeness* of the internal representations. The main objective of these experiments was to test whether sufficient representations of a changed item are available when observers report no awareness of a change. In what form these representations are held—implicit or explicit, verbal or visual, in a short- or long-term store—is theoretically important and worthy of further investigation, but it is secondary to the focus of this article. Regardless of, say, whether the underlying representations are verbal or visual, our conclusions remain the same; that is, observers sometimes fail to detect changes even when they have access to representations of both the pre- and postchange information. In other words, change blindness can result from a comparison failure. Furthermore, it is not clear whether the retained representations are complete or partial; observers may only succeed at detecting the changes and responding to the 2AFC questions when they have retained a perfect representation of the pre- and postchange objects, or they may be able to succeed when they have formed partial, yet sufficiently detailed, representations. The present experiments were not designed to differentiate these two alternatives, and it is important to note that a comparison failure explanation does not rely upon complete representations. Rather, the representations need only to be sufficiently detailed to support change detection.

Conclusion

These four experiments provide evidence for access to preserved representations of the pre- and postchange information in the face of change blindness, suggesting that change blindness does not imply a lack of internal representations of the outside world. Change blindness neither logically (e.g., Simons, 2000) nor empirically requires the absence of internal representations. Not only do we form representations, but we form multiple representations that can be used to make multiple discriminations. The representations might be somewhat fragile and easily overwritten or disrupted, but they are sufficiently long-lived to allow for successful recognition performance. Although change detection appears to be a unitary act, it could fail for a number of reasons, any of which would cause change blindness. Our results add direct evidence for the critical importance of a successful comparison process in change detection.

REFERENCES

Angelone, B. L., Levin, D. T., & Simons, D. J. (2003). The roles of representation and comparison failures in change blindness. *Perception*, **32**, 947-962.

Beck, M. R., & Levin, D. T. (2003). The role of representational volatility in recognizing pre- and postchange objects. *Perception & Psychophysics*, **65**, 458-468.

Becker, M. W., Pashler, H., & Antis, S. M. (2000). The role of iconic memory in change-detection tasks. *Perception*, **29**, 273-286.

Brawn, P. T., Snowden, R. J., & Wolfe, J. M. (1999). The minimal conditions for "change blindness": What is hides what was [Abstract]. *Investigations in Ophthalmology & Visual Science*, **40**, S49.

Gibson, J. J. (1966). The problem of temporal order in stimulation and perception. *Journal of Psychology*, **62**, 141-149.

Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.), *Perception* (Vol. 2, pp. 89-110). New York: Oxford University Press.

Hayhoe, M. (2000). Vision using routines: A functional account of vision. *Visual Cognition*, **7**, 43-64.

Hollingworth, A. (2003). Failures of retrieval and comparison constrain change detection in natural scenes. *Journal of Experimental Psychology: Human Perception & Performance*, **29**, 388-403.

Hollingworth, A., & Henderson, J. M. (2002). Accurate visual memory for previously attended objects in natural scenes. *Journal of Experimental Psychology: Human Perception & Performance*, **28**, 113-136.

Hollingworth, A., Schrock, G., & Henderson, J. M. (2001). Change detection in the flicker paradigm: The role of fixation position within the scene. *Memory & Cognition*, **29**, 296-304.

Hollingworth, A., Williams, C. C., & Henderson, J. M. (2001). To see and remember: Visually specific information is retained in memory from previously attended objects in natural scenes. *Psychonomic Bulletin & Review*, **8**, 761-768.

Irwin, D. E. (1992). Memory for position and identity across eye movements. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 307-317.

Levin, D. T., Simons, D. J., Angelone, B. L., & Chabris, C. F. (2002). Memory for centrally attended changing objects in an incidental real-world change detection paradigm. *British Journal of Psychology*, **93**, 289-302.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.

McConkie, G. W., & Rayner, K. (1976). Identifying the span of effective stimulus in reading: Literature review and theories of reading. In H. Singer & R. B. Ruddell (Eds.), *Theoretical models and processing of reading* (pp. 137-162). Newark, DE: International Reading Association.

Noë, A., Pessoa, L., & Thompson, E. (2000). Beyond the grand illusion: What change blindness really teaches us about vision. *Visual Cognition*, **7**, 93-106.

O'Regan, J. K. (1992). Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, **46**, 461-488.

O'Regan, J. K., & Noë, A. (2002). A sensorimotor account of vision and visual consciousness. *Behavioral & Brain Sciences*, **24**, 939-1031.

Pashler, H. (1988). Familiarity and visual change detection. *Perception & Psychophysics*, **44**, 369-378.

Phillips, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics*, **16**, 283-290.

Poggio, T., Torre, V., & Koch, C. (1985). Computational vision and regularization theory. *Nature*, **317**, 314-319.

Rensink, R. A. (2000). The dynamic representation of scenes. *Visual Cognition*, **7**, 17-42.

Rensink, R. A. (2002). Change detection. *Annual Review of Psychology*, **53**, 245-277.

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, **8**, 368-373.

Rosenthal, R., & Rosnow, R. L. (1991). *Essentials of behavioral research: Methods and data analysis* (2nd ed.). New York: McGraw-Hill.

Ryan, J. D., & Cohen, N. J. (2004). The nature of change detection and online representation of scenes. *Journal of Experimental Psychology: Human Perception & Performance*, **30**, 988-1015.
Scott-Brown, K. C., Baker, M. R., & Orbach, H. S. (2000). Comparison blindness. *Visual Cognition*, **7**, 253-267.
Shore, D., & Klein, R. M. (2000). The effects of scene inversion on change blindness. *Journal of General Psychology*, **127**, 27-43.
Silverman, M., & Mack, A. (2001). Priming from change blindness [Abstract]. *Journal of Vision*, **1** (3), 13a.
Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition*, **7**, 1-15.
Simons, D. J., Chabris, C. F., Schnur, T., & Levin, D. T. (2002). Evidence for preserved representations in change blindness. *Consciousness & Cognition*, **11**, 78-97.
Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning & Memory*, **6**, 174-215.
Stroud, J. M. (1955). The fine structure of psychological time. In H. Quastler (Ed.), *Information theory in psychology: Problems and methods* (pp. 174-207). Glencoe, IL: Free Press.
Tatler, B. W. (2001). Characterising the visual buffer: Real-world evidence for overwriting early in each fixation. *Perception*, **30**, 993-1006.
Trehub, A. (1991). *The cognitive brain*. Cambridge, MA: MIT Press.
Wolfe, J. M. (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli* (pp. 71-94). Cambridge, MA: MIT Press.

NOTES

1. All p values associated with t tests are two-tailed. All t tests associated with A', B'', or correlational values are one-sample t tests comparing the value with chance performance (e.g., $A' = 0.5$).

2. One concern is that the frequency with which observers report detecting a change directly affects the number of trials contributing to their accuracy scores for the aware and unaware conditions. That is, when observers detect more changes, they contribute fewer unaware trials. However, the unequal number of trials in these conditions had no effect on the overall values; weighted means that took into account the number of contributing trials did not differ from the unweighted means. All values reported in this article are unweighted.

3. All mean correlations reported in this article were calculated by converting each observer's correlation to Fisher's r_z to place it in a normal distribution, taking the mean, and then converting the mean back to the corresponding value of r (Rosenthal & Rosnow, 1991).

4. One additional analysis provides further support for the similarity of the representations underlying change detection and recognition performance; observers rarely were incorrect on both the pre- and post-change 2AFC questions when they reported being aware of the change (<2.5% of the aware trials across Experiments 1–3). Such a finding is expected if the representations for the two tasks are related.

APPENDIX A
2AFC Question Response Times

In the table below, response times (RTs) and standard deviations (*SD*s) are provided for both aware and unaware trials (all values in milliseconds).

Table A1

	Aware		Unaware	
Question	RT	*SD*	RT	*SD*
	Experiment 1			
Prechange	972.93	133.64	1,152.41	242.94
Postchange	786.86	102.01	974.76	187.31
Unchanged	1,079.59	174.96	1,007.61	164.10
	Experiment 2			
Prechange	923.55	188.56	1,172.96	204.68
Postchange	659.34	108.36	974.59	219.10
	Experiment 3			
Prechange	978.42	230.55	1,092.47	236.46
Postchange	738.58	227.87	1,002.60	262.37

APPENDIX B
2AFC Question Presentation Order Effects

Given that the 2AFC questions were asked in a random order on each trial, accuracy might have been affected by the added time and potential for proactive interference from the first to the last question. However, as seen in the tables below, the presentation order did not affect recognition performance. (In the tables, values are collapsed across awareness.)

Experiment 1
Mean Accuracy by Order of Presentation for the Prechange, Postchange, and Unchanged Object Questions

	Presentation						
	First		Second		Third		
Question	*M*	*SD*	*M*	*SD*	*M*	*SD*	*F* Test
Prechange	72.53	8.80	69.30	8.45	70.64	7.93	$F(2,30) = 1.18$, $p = .320$
Postchange	85.24	6.05	85.30	7.71	83.45	7.02	$F(2,30) = 1.16$, $p = .328$
Unchanged	75.01	7.16	72.33	7.42	69.99	7.32	$F(2,30) = 7.40$, $p = .002$

Experiment 2
Mean Accuracy by Order of Presentation and Awareness for the Pre- and Postchange Object Questions

		Presentation				
		First		Second		
Question	Awareness	*M*	*SD*	*M*	*SD*	*t* Test
Prechange	Aware	74.31	8.00	77.60	8.31	$t(11) = 1.25$, $p = .238$
	Unaware	61.36	9.09	64.26	7.74	$t(11) = 1.20$, $p = .254$
Postchange	Aware	94.12	2.97	95.85	4.09	$t(11) = 1.64$, $p = .130$
	Unaware	78.72	5.00	73.00	10.81	$t(11) = 2.18$, $p = .052$

Experiment 3
Mean Accuracy by Order of Presentation and Awareness for the Pre- and Postchange Object Questions

		Presentation				
		First		Second		
Question	Awareness	*M*	*SD*	*M*	*SD*	*t* Test
Prechange	Aware	74.20	8.22	74.49	9.26	$t(11) = 0.12$, $p = .905$
	Unaware	54.76	10.66	56.85	11.39	$t(11) = 0.49$, $p = .635$
Postchange	Aware	93.06	5.07	91.47	7.14	$t(11) = 1.11$, $p = .293$
	Unaware	76.74	12.85	71.22	8.68	$t(11) = 1.83$, $p = .094$

(Manuscript received January 9, 2003;
revision accepted for publication November 20, 2003.)