Psychon Bull Rev (2015) 22:722–727
DOI 10.3758/s13423-014-0745-8

BRIEF REPORT

Inattentional blindness reflects limitations on perception, not memory: Evidence from repeated failures of awareness

Emily J. Ward · Brian J. Scholl

Published online: 17 December 2014

Abstract Perhaps the most striking phenomenon of visual awareness is inattentional blindness (IB), in which a surprisingly salient event right in front of you may go completely unseen when unattended. Does IB reflect a failure of perception, or only of subsequent memory? Previous work has been unable to answer this question, due to a seemingly intractable dilemma: ruling out memory requires immediate perceptual reports, but soliciting such reports fuels an expectation that eliminates IB. Here we introduce a way of evoking repeated IB in the same subjects and the same session: we show that observers fail to report seeing salient events' not only when they have *no* expectation, but also when they have the *wrong* expectations about the events nature. This occurs when observers must immediately report seeing anything unexpected, even mid-event. Repeated IB thus demonstrates that IB is aptly named: it reflects a genuine deficit in moment-by-moment conscious perception, rather than a form of inattentional amnesia.

Keywords Inattentional blindness · Visual awareness · Iconic memory

Introduction

Inattentional blindness (IB) — the failure to consciously perceive otherwise salient events when they are not attended (Mack & Rock, 1998; Most et al., 2005; Neisser & Becklen, 1975) — is one of psychology's biggest exports. The failure to perceive a gorilla walking through a scene, for example, has become entrenched in popular culture (Chabris & Simons, 2010), and its initial demonstration has been cited over a thousand times (Simons & Chabris, 1999). At least three factors may contribute to this impact: First, the magnitude of IB is often extreme: what we fail to perceive is not a subtle detail of an event, or a change from one feature to another, but rather the entire event itself, no matter how salient (e.g., a bright red cross moving through a field of black and white shapes; Most et al., 2001). Second, it is deeply counterintuitive: almost everyone is surprised that failures of awareness of this magnitude are even possible. And third, IB has dramatic real-world consequences (e.g., causing traffic accidents, or errors in medical diagnosis; e.g., Drew et al., 2013).

E. J. Ward (✉) · B. J. Scholl (✉)
Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA
e-mail: emily.ward@yale.edu
e-mail: brian.scholl@yale.edu

Blindness or amnesia?

What is the underlying cause of IB? There are at least two possibilities. One is that IB may reflect a failure of perception: we may fail to report seeing a gorilla because we never saw it in the first place. Alternatively, since such reports are solicited after the fact, IB may reflect a failure of memory: we may be phenomenally aware of 'missed' events, but inattention may nevertheless prevent encoding into memory — yielding not 'blindness,' but rather a form of 'inattentional amnesia' (e.g., Wolfe, 1999) or 'inattentional inaccessibility' (e.g., Block, 2011). Critically, note that intuition cannot settle this question. It certainly seems odd to think that we could see a gorilla but later be unable to report it, but of course such a phenomenon would by definition not be noticeable, and so could not influence our intuitions.

Though this issue seems foundational, previous work has been unable to resolve it, due to a seemingly intractable dilemma. If observers are asked about their perception of unexpected events after the fact, then it may simply be too late; perhaps they saw the events, but the resulting iconic traces have disappeared (cf., Sperling, 1960), with no encoding into more durable forms of memory. But if observers are given proactive instructions to report unexpected events, then they attend to those events (having now been looking for them), and so IB is attenuated or eliminated. This has struck several theorists as an unsolvable problem — e.g., that "there are serious problems with any experimental effort to directly ask subjects if something is consciously perceived without attention", and that this "proves to be impossible because the demand to report on [an unexpected event] directs attention to [it]" (Wolfe, 1999).

Here we introduce a way to escape this dilemma, reporting two experiments that evoke repeated IB in the same observers, in the same session — where the unexpected event not only occurs multiple times (e.g., as in Mack & Clarke, 2012), but where IB occurs even when observers are asked about what they noticed after each instance. In particular, we show that IB occurs not only when observers have *no* expectation for the relevant events, but also when they have the *wrong* expectations about the relevant events' nature. Moreover, this occurs when observers must immediately report seeing anything unexpected, even mid-event.

Experiment 1: Repeated Inattentional Blindness

We employed a sustained inattentional blindness task (Most et al., 2001, 2005): observers viewed moving black and white L and T shapes, counting the number of times that a subset crossed the display's midline. On the fourth trial, an unexpected gray cross (the Unexpected Event; UE) appeared and traversed the entire display, and observers subsequently reported whether they had noticed it. After this event, observers were explicitly told to watch for subsequent UEs — and during some later trials, the gray cross appeared again. We theorized that this would induce expectation that was specific to the gray cross. The critical trial then occurred after three presentations of the gray cross (interspersed with additional no-UE trials), when a new 'unexpected' object traversed the display: a black 'E' shape. (These events are depicted in Fig. 1, though using the colors and numbers of shapes from Experiment 2.) We assessed observers' awareness of this second UE immediately after its offset.

Method

Participants

Observers were recruited and run online via Amazon Mechanical-Turk. (For discussion of this pool's reliability, see Crump et al., 2013.) Data collection continued until we had recruited 100 observers who failed to notice the first unexpected event (UE1), per the exclusion criteria discussed below (see Table 1).[1]

Apparatus

After agreeing to participate, observers were redirected to a web server where platform-independent stimulus presentation and data collection were completed by custom software run in observers' web browsers, written using a combination of PHP and jQuery.

Stimuli and procedure

Because the display loaded within observers' own web browsers, viewing distance and screen resolutions could vary dramatically, and so we report critical dimensions using absolute pixel (px) values. All events took place within a gray window (610 px ×500 px), horizontally bisected by a purple (1 px) line, and with a small (10 px) fixation square at its center. After 1 s, four black and four white L and T shapes (65 px block letters) appeared and moved independently for 14 s along linear paths at velocities randomly ranging from 144 to 229 px/s. During their motion, the shapes could occlude each other and reflect off the edges of the display, and passed under the midline. Observers were instructed to fixate the central point and to count how many times the white shapes crossed the midline. Following each trial, observers were prompted for their tally.

The initial trial sequence was typical for IB studies (e.g., Mack & Rock, 1998; Most et al., 2001), and began with three trials containing only the L and T shapes. Five seconds into the fourth trial, a gray cross (65 px; 37 % luminance) unexpectedly entered the display on the right, moved horizontally along the midline, passed behind the fixation point, and eventually exited the display on the left, after having been visible for 5 s.

After the motion ended, the display disappeared, and observers were asked whether they noticed "anything … that was different from the first three trials" — and if so, to describe what was different. They were then shown the gray cross and asked if they had noticed it — and if so, to describe where it was and how it moved. Only observers who explicitly reported not noticing the cross were counted as 'nonnoticers' to be included in the final sample ($N = 100$).

[1] Most previous IB experiments have used roughly 20 subjects per condition, but we arbitrarily decided before collection began to quintuple this number. We chose this relatively large sample not because it was necessary for our effects, but simply because online data collection allows for vastly larger samples to be readily collected. This is to our knowledge the first report of IB with online data collection.

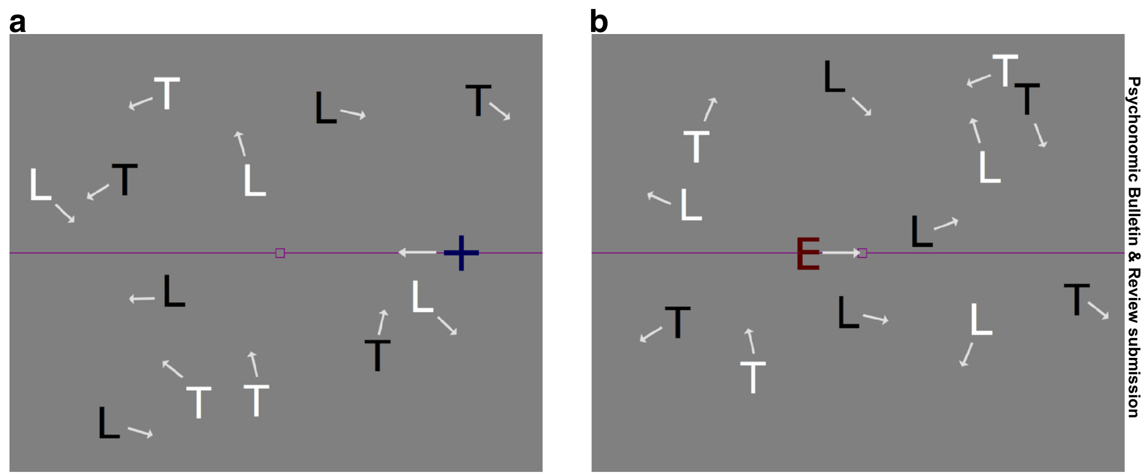

Fig. 1 (**a**) A sample screenshot of a sustained inattentional blindness display from Experiment 2 wherein the unexpected event is a dark blue cross traversing the display from right to left. (**b**) A sample screenshot of a sustained inattentional blindness display from Experiment 2 wherein the unexpected event is a dark red 'E' shape traversing the display from left to right. In Experiment 2, one of these shape/motion combinations (counterbalanced across observers, with color also counterbalanced) was used for UE1, UE2, and UE3, while the other shape/motion combination was used for UE4. In Experiment 1, UE1, and UE2, and UE3 were always gray crosses, and UE4 was a black E

Observers then completed six more trials. Trials five, seven, and eight contained no UE; trials six (UE2) and nine (UE3) contained the gray cross again, as 'divided attention' trials (e.g., Most et al., 2001) — and after each of these trials (i.e., every trial after UE1), observers were asked the same series of questions that had followed UE1. On the tenth and last trial (UE4), a black E shape (65 px) unexpectedly entered the display on the left after 5 s, moved horizontally along the midline, passed behind the fixation point, and eventually exited the display on the right, after having been visible for 5 s. Immediately after UE4 disappeared, the display disappeared, and observers were again asked the same series of questions (though now probing about the black E rather than the gray cross). Observers who explicitly reported not noticing the black E were counted as 'repeated nonnoticers.'

After the 10th trial ended, observers answered follow-up questions about their performance on the midline-crossings task, their attention to instructions, their expectations about unexpected events, what they were looking for, and whether they had heard of similar experiments (e.g., with a missed gorilla). We excluded from analyses observers who left questions unanswered or failed to answer an instruction check, and those who reported expecting an unexpected event. These exclusions were computed during the data collection phase, so that the experiment resulted in exactly 100 nonnoticers after these exclusions (see Table 1 for details).

UE4 differed from the previous UEs not only in its use of a different shape and motion direction, but also in that it came later in the session. To control for this, we also included a between-subjects condition in which the critical final trial (UE4) again

Table 1 Incidence of repeated inattentional blindness. The number of observers recruited and included in the final samples for Experiments 1 and 2, along with the resulting degrees of inattentional blindness for the initial UE (UE1) and the critical UE (UE4) are presented. In both experiments, data collection continued until exactly 100 non-excluded observers who missed UE1

Critical Unexpected Event Type	Expt. 1 – Repeated IB		Expt. 2 – Immediate Report	
	Black E	Gray Cross (Control)	Novel Event	Repeated Event (Control)
Total Recruited	**239**	**302**	**468**	**565**
Total Excluded	**42**	**67**	**135**	**168**
Colorblind	-	-	46	42
Expecting unexpected	40	59	74	97
Poor/lack of response	2	8	15	29
Total Included	**197**	**235**	**333**	**397**
# who missed UE1	**100 (50.76 %)**	**100 (42.55 %)**	**100 (30.03 %)**	**100 (25.19 %)**
# who missed UE4	**29 (29.00 %)**	**16 (16.00 %)**	**13 (13.00 %)**	**4 (4.00 %)**

The bold text is to differentiate the key data from the exclusion criteria consisted of the gray cross appearing for its fourth time. This allowed us to directly compare the number of observers who failed to notice the novel type of UE4 with the repeated type of UE4 — and so ensure that any differences could not simply reflect its serial position within the trial sequence.

Results

Error rates for the overt midline-crossings task (i.e., the difference between their tally and actual number, divided by the actual number) were not significantly different for observers who noticed (22.67 %) or failed to notice (18.75 %) the UE1 ($t(195) = 1.52$, $p = 0.13$, $d = 0.22$).

As shown in Table 1, approximately half of observers failed to notice UE1 (100/197 = 50.76 %), replicating the phenomenon of IB. (Note that IB rate varies as a function of many factors, such as the visual similarity of the UE to the task-relevant items; e.g., Most, et al. 2001.)

Critically, of these 100 nonnoticers, 29 subsequently missed the final unexpected event when it had unexpected features (i.e., the black rightward-moving E), even when they were asked to report their experience immediately after its offset (during which time a report could still be based on the visual 'icon,' had one formed; see Wolfe, 1999). To our knowledge, this is the first demonstration of *repeated inattentional blindness* within the same individual, in the same session, with awareness assessed after each instance.[2]

Importantly, this repeated IB was not simply due to the position of the UE in the trial sequence: only 16 of 100 observers in the time-matched control experiment failed to see the repeated UE4 (i.e., the previously displayed leftward-moving gray cross) — which was significantly fewer than in the experimental group ($\chi^2(1) = 4.85$, $p = 0.03$; $\phi = 0.16$). That the repeated IB was due to novel properties of UE4 can also be seen by looking at what happened on UE2 and UE3 in the primary versus control experiments. Of observers who missed UE4 when it was a black E, 36.84 % also missed either UE2 or UE3; of observers who missed UE4 when it was a repeated gray cross, 44.44 % also missed either UE2 or UE3. Thus the groups were equally attentive ($\chi^2(1) = 0.27$, $p = 0.6$; $\phi = 0.08$), yet almost twice as many observers missed UE4 when it had novel features compared to when it had repeated features.

Experiment 2: Repeated IB with Immediate Perceptual Reports

Inspired by the fact that repeated IB was possible at all, we next attempted to assess it with instructions to immediately report UEs, even mid-trial — with no time for perceptual decay. After the first UE, observers pressed a key any time throughout the rest of the experiment when they saw something different or unexpected. As in Experiment 1, each observer saw the same UE repeated multiple times (as UE1, UE2 and UE3), followed by a novel UE4. Of interest was whether they would indicate seeing the novel UE4 as it was happening, at a greater rate than a repeated UE4. In addition, to ensure that any such difference was not due to brute visual differences between the UE types, we counterbalanced which stimuli were used for the repeated versus novel UEs across observers — also now making the UEs more salient by presenting them in novel colors (different from every other object in the display; see Fig. 1).

Method

This experiment was identical to Experiment 1 except as reported here. To increase the salience of the UEs — and to make them equally distinct from the task-relevant items and from each other — the cross and the E shapes were now colored (see Fig. 1), dark red (RGB 90/0/0) or dark blue (RGB 0/0/90), with color counterbalanced across the groups described below. The number and speed of the moving shapes was increased, to six black and six white L and T shapes, with velocities ranging from 170 to 255 px/s; this increased the difficulty of the primary task, and compensated for the increased salience of the UEs. A new group of observers was run until we had obtained 100 observers ('nonnoticers') who missed UE1, using the same exclusion criteria, but now also excluding colorblind observers, as assessed by an incorrect response to an Ishihara plate presented during the follow-up questions (see Table 1). Half of the observers saw the leftward-moving cross as UE1-UE3, and the rightward-moving E as UE4 (as in Experiment 1); the other half saw the rightward-moving E as UE1-UE3, and the leftward-moving cross as UE4. In the corresponding control group, 50 observers saw just the cross and 50 observers saw just the E.

After the trial with UE1 had ended, observers were told that: "any time through the rest of the experiment that you see something different or unexpected …, immediately press the space bar and we'll ask you to identify it." After that point, a

[2] We thank Steve Most for indirectly inspiring this manipulation. Once, while testing observers in a prior IB study, Most and the second author of the current report were both testing observers in the same room, two at a time. Normally, as an experimenter in IB studies, one may stand behind the subject, gobsmacked that the observers fail to see the UE (here a red cross) — which is so obvious to you. But once, the second author left to use the bathroom mid-session, and on the subsequent two subjects after he returned, he also failed to see the red cross, despite actively looking for it! He then stopped the testing, explaining to Most that there must be some bug in the code that was keeping the UE from appearing — at which point Most gently explained to him that during his absence the red cross subjects had been finished, and that Most had set both computers to start running green-triangle UE trials. Armed with this mistaken expectation, though, the second author had seen nothing — thus experiencing IB for the first time himself.

given trial always ended immediately upon any keypress, after which observers were asked the same series of questions from Experiment 1 in order to identify what had led them to press the key. If no key was pressed during a trial, observers were still asked the usual series of questions at the end, as in Experiment 1.

Results

Error rates for the overt midline-crossings task were very similar to those of Experiment 1, and were not significantly different for observers who noticed (25.03 %) or failed to notice (27.29 %) the UE1 ($t(329) = 0.70$, $p = 0.49$, $d = 0.08$).

As shown in Table 1, approximately one-third of observers failed to notice UE1 (100/333 =30.03 %), again replicating the phenomenon of IB. (This is a lower proportion than in Experiment 1 [$\chi^2(1) = 22.54$, $p<0.001$; $\phi = 0.21$], presumably due to the increased salience of the now-colored UEs.) Critically, of these 100 nonnoticers, 13 subsequently missed the final unexpected event when it had unexpected features (i.e., the rightward-moving red E after previously encountering only leftward-moving blue crosses, or vice versa, with colors counterbalanced). And since the shapes and motions used for the two types of UEs were counterbalanced, this difference could not be due to the second UE just being less noticeable in general. To our knowledge, this is the first demonstration of repeated inattentional blindness — or of any IB at all! — when observers had the explicit task to immediately respond to anything unexpected.

Again, this repeated IB was not simply due to the position of the UE in the trial sequence: only four of 100 observers in the time-matched control experiment failed to see the repeated UE4 — which was significantly fewer than in the experimental group ($\chi^2(1) = 5.21$, $p = 0.02$; $\phi = 0.16$). Thus, even when given explicit instructions to monitor for and respond immediately to unexpected events, *more than three times* as many people miss a UE that does not match their expectations. And since the shapes and motions used for the two types of UEs were counterbalanced, this difference could not be due to simple visual differences (e.g., if the UE4 used in the novel condition was just less noticeable in general than the UE4 used for the repeated condition — which was a possibility in Experiment 1).

Of course, given the instruction to report only further 'unexpected' events, many observers (27 % in the experimental group, 23 % in the control group) did not press the key when further repeated events appeared after UE1 (after which point they were no longer 'unexpected' by their definitions) — even while they might still later notice the novel UE4. To ensure that this dynamic did not distort the results, we also measured IB in only the subset of observers who did hit the immediate-report/noticing key in response to at least one familiar event after UE1 but before the critical UE4 (i.e., in response to UE2 or UE3). This analysis qualitatively replicated the primary pattern of results: 12.33 % of these observers failed to notice the novel UE4, while only 1.30 % (i.e., nearly ten times fewer) failed to notice the repeated UE4 ($\chi^2(1) = 7.33$, $p = 0.01$; $\phi = 0.22$).

Discussion

As with Experiment 1, the key aspect of these results is not the brute magnitude of repeated IB (though a figure of 12–13 % would certainly still be tremendously ecologically relevant in the context of, say, traffic), but rather the (immediate report) conditions under which it occurred. And this final test — of IB rates for observers who used the immediate-response option — represents the strongest test of our hypothesis: repeated inattentional blindness was still experienced even by many observers who (a) were explicitly instructed to attend to unexpected events, (b) were able to immediately report (even mid-trial) anything unexpected, and (c) actively made use of this ability on previous trials.

General Discussion

The two experiments reported here have a straightforward methodological implication, and a straightforward theoretical implication.

Methodologically, these results demonstrate that it is possible to obtain inattentional blindness repeatedly in the same observers, in the same session, even while assessing awareness after each critical event. The key to this demonstration was that IB can arise not only from a lack of any expectation, but also from a *wrong* expectation (about what will occur) — where this difference increases the magnitude of IB at least threefold in some contexts (in Experiment 1), and possibly almost tenfold in others (in Experiment 2). This is the first time to our knowledge that repeated IB has been reported, where the second 'missed' event is of the same general character as the initially 'missed' event (here differing in shape, motion direction, and color). Previously, Simons (2010) demonstrated that even observers who noticed a not-entirely-unexpected gorilla could still subsequently miss other unexpected features — but these were very different display details, better characterized as change-blindness, e.g., a change in background curtain color. Failures to detect changes can also be striking, but not nearly as striking as the complete failure to see a novel UE at all: e.g., hitting a pedestrian while driving might be caused by a failure to see her at all, but not by a failure to notice that her shirt-color changed.

Theoretically, these results address a foundational issue that other theorists have worried was experimentally intractable.[3] Because previous demonstrations of IB always assessed noticing after the UE had disappeared, they could always be reinterpreted as (mere) failures to encode visual information into memory, despite a possibly rich 'in the-moment' phenomenological experience (e.g., Block, 2011; Wolfe, 1999). (Similarly, any demonstration of change blindness can be explained by memory/comparison failures rather than a lack of visual experience; e.g., Mitroff et al., 2004). But the current results show that IB persists even when observers report the UE during online perception — with no time for perceptual decay whatsoever (indeed, responding while the UE is fully visible). We thus conclude that IB is aptly named: it reflects a genuine deficit in moment-by-moment conscious perception, rather than inattentional amnesia or inaccessibility.

Acknowledgments For helpful conversation, inspiration, and/or comments on earlier drafts, we thank San-Yuan Lin, Gary Lupyan, Steve Most, Dan Simons, Jeremy Wolfe, and Su-Ling Yeh.

References

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences, 15,* 567–575.

Chabris, C. F., & Simons, D. J. (2010). *The invisible gorilla: How our intuitions deceive us*. New York: Random House.

Crump, M., McDonnell, J., & Gureckis, T. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PloS One, 8,* e57410.

Drew, T., Vo, M., & Wolfe, J. (2013). The invisible gorilla strikes again: Sustained inattentional blindness in expert observers. *Psychological Science, 24,* 1848–1853.

Mack, A., & Clarke, J. (2012). Gist perception requires attention. *Visual Cognition, 20,* 300–327.

Mack, A., & Rock, I. (1998). *Inattentional blindness*. Cambridge, MA: MIT Press.

Mitroff, S., Simons, D., & Levin, D. (2004). Nothing compares two views: Change blindness can occur despite preserved access to the changed information. *Perception & Psychophysics, 66,* 1268–1281.

Most, S. B., Scholl, B. J., Clifford, E., & Simons, D. J. (2005). What you see is what you set: Sustained inattentional blindness and the capture of awareness. *Psychological Review, 112,* 217–242.

Most, S. B., Simons, D. J., Scholl, B. J., Jimenez, R., Clifford, E., & Chabris, C. F. (2001). How not to be seen: The contribution of similarity and selective ignoring to sustained inattentional blindness. *Psychological Science, 12,* 9–17.

Neisser, U., & Becklen, R. (1975). Selective looking: Attending to visually specified events. *Cognitive Psychology, 7,* 480–494.

Rees, G., Russell, C., Frith, C. D., & Driver, J. (1999). Inattentional blindness versus inattentional amnesia for fixated but ignored words. *Science, 286,* 2504–2507.

Simons, D. J. (2010). Monkeying around with the gorillas in our midst: Familiarity with an inattentional blindness task does not improve the detection of unexpected events. *i-Perception, 1,* 3–6.

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception, 28,* 1059–1074.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied, 74,* 1–29.

Wolfe, J. M. (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli* (pp. 71–94). Cambridge, MA: MIT Press.

[3] Only one previous study to our knowledge attempted to distinguish inattentional blindness from amnesia (Rees et al., 1999), but we contend that it failed to do so and did not even truly test IB in the sense used here. Observers saw rapid streams of letters superimposed on pictures, and monitored one of the streams. Using functional imaging, differential brain responses were found to real words compared to random letters when words were attended, but not when pictures were attended. This does not distinguish blindness from amnesia, since a lack of semantic processing does not imply a lack of visual content. More importantly, a failure to read words does not imply that observers did not *see* them. The authors of this study acknowledge this and "do not suggest that … participants were blind to the presence of the letters… but rather were blind to those properties that distinguish words from random strings of consonants." This may be an interesting finding, but it doesn't seem closely related to inattentional blindness, per se.