**Inattentional blindness on the full-attention trial:
Are we throwing out the baby with the bathwater?**

Rebekah C. White (1)*

Martin Davies (1, 2)

Anne M. Aimola Davies (1, 3)

(1) Department of Experimental Psychology, University of Oxford

(2) Faculty of Philosophy, University of Oxford

(3) Research School of Psychology, Australian National University

* All authors contributed equally to this work.

Address correspondence to: Rebekah.White@psy.ox.ac.uk; Martin.Davies@ccc.ox.ac.uk;

Anne.Aimola@anu.edu.au

Abstract

When attention is otherwise engaged, observers may experience *inattentional blindness*, failing to notice objects or events that are presented in plain sight. In an inattentional blindness experiment, an unexpected stimulus is presented alongside primary-task stimuli, and its detection is probed. We evaluate a criterion that is commonly used to exclude observers from the data analysis. On the final experimental trial, observers do not perform the primary task, but instead look for anything new. Observers who fail to report the unexpected stimulus on this *full-attention trial* are excluded. On the basis of 4 hypothetical experiments and a review of 125 actual experiments from the literature, we demonstrate some potentially problematic consequences of implementing the full-attention-trial exclusion criterion. Excluded observers may cluster in experimental conditions and the exclusion criterion may lead researchers to understate the pervasiveness of inattentional blindness. It may even render us *blind* to *inattentional blindness* on the *full-attention trial*.

Keywords: Attentional set, Data screening, Exclusion criterion, Forced-choice display, Full-attention trial, Inattentional blindness

1. Introduction

When an observer's attention is taken up with a task, she may fail to notice an unexpected stimulus (an object or event) that is presented in plain sight (Neisser & Becklen, 1975) and precisely where her attention is directed as she carries out her task. Real-world examples of inattentional blindness can be amusing (e.g., failing to notice a friend who is waving to capture one's attention; Simons, 2000) or they can be devastating (e.g., failing to notice an obstacle in the path of one's car; Most & Astur, 2007). Therefore considerable importance attaches to research that advances our understanding of this striking phenomenon.

In most inattentional blindness experiments, following trials in which the observer's attention is engaged by a task, there is a final *full-attention trial*, which is treated as a control trial. Observers who do not notice the unexpected stimulus even on this control trial are excluded from subsequent data analysis. This is the *full-attention-trial exclusion criterion*. We shall argue, first, that this exclusion criterion is potentially problematic. Some differences in reported rates of inattentional blindness in the literature (differences between experiments or differences between conditions within a single experiment) may reflect the exclusion criterion, rather than being explained by the experimenter's theoretically motivated manipulations.

Failures of noticing on the full-attention trial sometimes cluster in experimental conditions that produce high rates of inattentional blindness on the critical trial. We shall argue, second, that these failures of noticing may be evidence, not of the need to exclude observers, but of something theoretically more interesting – a form of genuine inattentional blindness on the full-attention trial (paradoxical though that may sound). Investigating how, and in which experimental paradigms and conditions, this continued failure to notice the unexpected stimulus arises, may enhance our understanding of inattentional blindness itself.

To demonstrate the potentially problematic implications of implementing the full-attention-trial exclusion criterion, we shall present four hypothetical experiments (Section 4). To support the suggestion that some failures of noticing on the full-attention trial result from continuing inattentional blindness, we shall provide an account of how this might arise (Section 5). We have also compiled a table with 125 actual experiments (56 studies) from the inattentional blindness literature, detailing the structure of the trials used in these studies and the percentage of observers excluded from each experiment on the basis of the full-attention trial. The table includes, for each experiment, an indicative maximum percentage of observers

excluded from a condition if (hypothetically) all the observers excluded from the experiment clustered in a single condition (Section 3).

We outline the potential problems with the full-attention-trial exclusion criterion, and argue for the possibility of inattentional blindness on the full-attention trial, in the hope that this will stimulate further discussion among researchers. For readers who are less familiar with inattentional blindness, we preface our discussion with a brief review of the methods used for investigating inattentional blindness.

2. *Inattentional blindness experiments*

2.1. *Static and dynamic paradigms*

Static inattentional blindness paradigms involve primary tasks with brief presentation times, ranging from approximately 200 ms (Mack & Rock, 1998) to 1000 ms (Koivisto & Revonsuo, 2007). The observer performs a detection or discrimination task. In Mack and Rock's (1998) pioneering inattentional blindness experiments, the primary task was to discriminate the longer arm of a briefly-presented cross. The unexpected stimulus was a simple geometric shape (e.g., a small square), which was presented alongside the cross for 200 ms. Depending on the location of the unexpected geometric shape in the display, it went undetected by as many as 80% of observers (see also Mack and Rock, 1999).

Dynamic inattentional blindness paradigms involve primary tasks with longer presentation times, ranging from approximately 8 seconds (Simons & Jensen, 2009) to 75 seconds (Simons & Chabris, 1999). The observer performs a visual-tracking task (see Most et al., 2001; 2005, Simons & Chabris, 1999). In what is arguably the best-known inattentional blindness study (Simons & Chabris, see Neisser and Becklen, 1975 who inspired the paradigm), observers watched a video depicting two teams of basketball players passing a ball. The two teams could be distinguished from one another by the colour of their t-shirts – white or black. The task was to count the number of passes made by one of the two teams, as specified by the experimenter. In the easy version of the task, the observer kept a single count of passes, including bounce passes and aerial passes, and in the hard version of the task, the observer kept separate counts of bounce passes and of aerial passes. In one version of the video, the unexpected event was a person, dressed in a gorilla suit, walking through the centre of the action, remaining clearly visible for 5 seconds. Rates of inattentional blindness varied, depending on which team's passes were to be counted and on the level of difficulty of the counting task. When observers were doing the easy version of the task, the 'gorilla' went undetected by 58% of observers attending to the team wearing

white t-shirts and by 17% of observers attending to the team wearing black t-shirts. For observers attending to the team wearing black t-shirts and doing the hard version of the task, the rate of inattention blindness increased from 17% to 42%. This latter difference is consistent with explanations for inattention blindness based on the demands of the primary task (e.g., Cartwright-Finch & Lavie, 2007; White & Aimola Davies, 2008).

Most et al. (2001; see also 2005) incorporated many features of Simons and Chabris' (1999) experiment in a computer-based dynamic paradigm. Black and white Ls and Ts moved around a display (with a grey background), occasionally bouncing off the edges of the display window. The primary task was to count the bounces made by the letters of one colour (either black or white) during a 15-s trial. Letters of the other colour were *distractors* and could be ignored. The unexpected stimulus was a cross (the same size as the letters), which moved across the display, remaining visible for 5 s. Of 16 observers whose task was to count the bounces by the *white* letters, 15 (94%) failed to notice a *black* cross as it moved across the display. In contrast, only one of 16 observers (6%) whose task was to count the bounces by the *black* letters failed to notice the *black* cross. When primary-task stimuli include distractors as well as targets, observers are much less likely to notice an unexpected stimulus that is similar to the distractors (and so is *incongruent* with the observers' attentional set) than an unexpected stimulus that is similar to the targets (and so is *congruent* with the observers' attentional set).

2.2. *The unexpected stimulus: Critical trial, divided-attention trial, full-attention trial*

In some studies, there is only one trial of the primary task, and the unexpected stimulus is presented on this trial (e.g., the gorilla: Simons & Chabris, 1999). But in the majority of studies, there are at least two (pre-critical) trials of the primary task before the unexpected stimulus is presented. Following these pre-critical trials, the unexpected stimulus may be presented once (i.e., on the critical trial), twice (i.e., on the critical and full-attention trials), or three times (i.e., on the critical, divided-attention, and full-attention trials).

The first trial with an unexpected stimulus is referred to as the *critical trial* or the critical-inattention trial. Following this trial, the experimenter assesses whether the unexpected stimulus was noticed, by asking the observer a standard set of questions. For example, in Mack and Rock's (1998) experiment in which the primary task was to discriminate the longer arm of a briefly-presented cross, observers were asked "whether they had seen anything on the screen other than the cross figure, that is, anything that had not been present on previous trials" (p. 6). Similarly, in the study by Most et al. (2001), in which the primary task was to count the bounces by black or white Ls and Ts, observers were asked: "On the last trial, did you see anything other

than the black and white Ls and Ts (anything that had not been present on the first two trials)? If you did see something on the last trial that had not been present during the first two trials, please describe it” (p. 17).

Following these questions, the experimenter may present a *forced-choice* display and ask the observer to indicate which item in the display was present on the critical trial. For example, in one of Mack and Rock’s (1998) experiments, the unexpected stimulus was a cartoon happy face, sad face, or scrambled face, or a circular outline. If the unexpected stimulus was, for example, a sad face then, following the initial questions, the observer was shown a forced-choice display of five figures: a diamond, a sad face, a circle, a square, and a scrambled face. The observer was asked to select or guess the figure that was the unexpected stimulus (p. 142). In the study by Most et al. (2001), if the unexpected stimulus was a black cross, the observer viewed a forced-choice display containing four black geometric shapes: a capital E, a heart, a triangle, and a cross. The observer was asked: “If you did see something during the last trial that had not been present during the first two trials, please circle the shape of the object below. If you did not see anything, please guess. (Please indicate whether you did see something or are guessing)” (p. 17).

The second trial with an unexpected stimulus is referred to as the *divided-attention trial*. As a result of the experimenter’s questions following the critical trial, the observer has reason to expect that there may be something in addition to the primary-task stimuli, and may therefore divide attention between performing the primary task and looking for an additional item. After this trial, the experimenter once again assesses whether the unexpected stimulus was noticed, with the same standard set of questions as were asked after the critical trial. Following these questions, the experimenter *may* once again present a forced-choice display (e.g., Most et al., 2005) and ask the observer to indicate which item was present on the divided-attention trial.

The third and final trial with an unexpected stimulus is referred to as the *full-attention trial*. Immediately before this trial, the observer is told that he or she no longer needs to perform the primary task. The observer may also be explicitly instructed to look for any new item in the display. In Mack and Rock’s (1998) experiments, the observer was told, for example, to “ignore the cross [rather than judging which arm was longer], and to report only what else they saw on the screen” (p. 8). In the study by Most et al. (2001), the observer was told: “On this trial, the instructions are slightly different. As before, keep your eyes fixated on the fixation point, but this time don’t count the bounces any of the shapes make. Simply watch the display” (p. 11).

In Mack and Rock’s (1998) pioneering experiments, the crucial comparison was between noticing of the unexpected stimulus on the critical trial (when attention was otherwise engaged) and on the full-attention trial (when attention was available). On the critical trial, the observer’s

attention was engaged by the primary task and the observer had no reason to expect or search for an additional stimulus. The observer's answer to the experimenter's question following the critical trial "provides the data about what is perceived *without attention*" (p. 7; emphasis added). On the full-attention trial, the observer was instructed to ignore the primary task. So attention was potentially available to be directed to an additional stimulus and, as a result of the questions asked following the critical trial and the divided-attention trial, the additional stimulus was no longer unexpected. While noticing the unexpected stimulus on the critical trial would be noticing *without* attention, noticing it on the full-attention trial would be noticing *with* attention.

Thus, the difference in noticing performance between the two trials provided an indication of the contribution that attention (and expectation) made to noticing. Equivalently, the difference provided an indication of the contribution that inattention (and lack of expectation) made to an observer's 'blindness' – failure to notice the unexpected stimulus. (Mack and Rock (1998) discussed the role of expectation in chapter 9.) The intermediate divided-attention trial provided evidence about the degree to which observers could deploy attention so as to perform the primary task *and* notice the not-wholly-unexpected additional stimulus. Testing some 5,000 observers in hundreds of experimental conditions at two laboratories over a period of seven years, Mack and Rock established that inattention blindness is a robust phenomenon.

2.3. *The role of the full attention trial and the rise of the full-attention-trial exclusion criterion*

In order to establish that there is such a phenomenon as inattention blindness, Mack and Rock (1998) needed a clear contrast between noticing an unexpected stimulus under conditions in which attention was fully available and failure to notice the same stimulus under conditions in which attention was otherwise engaged. If an unexpected stimulus were to be reliably detected and identified on the critical (inattention) trial, just as on the final full-attention trial, there would be no evidence of inattention blindness. Rather, it would be clear that "attention is not required for its perception" (p. 10). More precisely, it would be clear that *voluntary* (endogenous) attention was not required for the unexpected stimulus to be noticed. Mack and Rock (chapter 5, pp. 116–9) showed that a highly salient stimulus that *captured* attention, the observer's own name, was reliably detected and identified on the critical trial, under conditions in which about 50% of observers failed to notice other words – even high-frequency words.

If an unexpected stimulus were to go undetected on the full-attention trial, just as on the critical trial, there would, again, be no evidence of inattention blindness. Rather, it would be clear (or apparently clear) that the unexpected stimulus was *imperceptible*, even with attention. As

Mack and Rock (1998) explained: “The control [full-attention] trial was important in establishing the *perceptibility* of the critical [unexpected] stimulus under conditions of brief, masked, and often parafoveal (rather than foveal) presentation *when attention was allowed*” (p. 10; emphasis added). In the project of establishing the phenomenon of inattentional blindness, the role of the full-attention trial as a check on perceptibility was dialectically important. Inattentional blindness had to be distinguished from blindness toward the imperceptible and it could not simply be assumed that a small stimulus (e.g., subtending a visual angle of 0.15°), briefly exposed (e.g., for 200 ms) and then masked, would be noticed, provided that attention was available.

Thanks to Mack and Rock’s (1998) seminal contributions, inattentional blindness is now an established phenomenon. Also, whereas Mack and Rock used brief presentations that were generally static and masked, in subsequent research there has been a shift toward also investigating inattentional blindness with prolonged, dynamic and unmasked presentations. In such experimental paradigms, failures to detect an unexpected stimulus are unlikely to result from its outright imperceptibility. It is reasonable, for example, to assume that observers in Simons and Chabris’ (1999) ‘gorilla’ experiment *would*, if attention were available, notice a person dressed in a gorilla suit, clearly visible for 5 seconds, walking through the scene. With these differences, attributions of inattentional blindness are no longer based solely on a difference in noticing performance between the critical and full-attention trials.

Some inattentional blindness experiments (e.g., Simons & Chabris, 1999) do not include a full-attention trial but many still do, even when there is no issue about the perceptibility of the unexpected stimulus. Thus, there has been a shift in the *rationale* for the full-attention trial. For example, Most et al. (2001) included a full-attention trial “as a control to ensure that [individual observers] could understand and follow task instructions” (p. 11). Kreitz, Furley and Memmert (2016a) suggested that failure to notice the unexpected stimulus on the full-attention trial indicated that “the participant potentially did not follow task instructions or has basal visual problems” (p. 499). Remington, Cartwright-Finch and Lavie (2014) suggested that failures to detect the unexpected stimulus on the full-attention trial “may have resulted from undiagnosed reduced visual acuity” (p. 8). In short, in the majority of studies, the full-attention trial is not used to check that the unexpected stimulus is perceptible (when attention is available). Rather, perceptibility of the stimulus is assumed and the full-attention trial is used to check that individual observers (a) understand and follow instructions, so that they engage properly with the experiment and (b) do not have undiagnosed or unreported visual problems.

With this shift in rationale there has also been a shift in the *role* of the full-attention trial. In Mack and Rock’s (1998) experiments, “subjects virtually always succeeded in seeing and

correctly identifying the critical stimulus ... in the control [full-attention] trials” (p. 10). Nevertheless, in some experiments, rates of detection on the full-attention trial were far from perfect (e.g., 73%; p. 120). Importantly, Mack and Rock (1998) *did not exclude* from their analyses observers who failed to report the unexpected stimulus on the full-attention trial. Indeed, because Mack and Rock conceived the full-attention trial as a check on the perceptibility of the *stimulus*, it would not have been natural for them to exclude *individual observers*. In more recent work, perceptibility of the stimulus is assumed and the full-attention trial is conceived as a check on the properties of individual observers – specifically, on whether they meet conditions (a) and (b). It is natural to exclude from data analysis those observers who do not follow instructions or who have visual problems. Consequently, to the extent that failure to report the unexpected stimulus on the full-attention trial indicates that the observer has not followed instructions or has visual problems (or both), the *full-attention-trial exclusion criterion* seems appropriate. When observers who fail to report the unexpected stimulus on the full-attention trial are excluded from the data analysis, there is an adjustment in the reported rates of inattentional blindness on the *critical* and *divided-attention trials*. This adjustment is naturally *downward*, because it is unlikely that an observer who noticed the unexpected stimulus on the critical trial or the divided-attention trial would fail to notice it on the full-attention trial.

In this paper, we suggest that researchers should exercise some caution when implementing the full-attention-trial exclusion criterion. The criterion is potentially problematic, as we became aware when analysing some of our own data (Aimola Davies et al., in preparation).

2.4. *Possible problems with the full-attention-trial exclusion criterion*

In a series of experiments (Aimola Davies et al., in preparation), we have been investigating how the observer’s attentional set predicts noticing of an unexpected stimulus. Our experiments use a static paradigm and the primary-task stimuli include distractors as well as targets. The unexpected stimulus replaces the fixation cross, and thus appears exactly where the observer is looking, and it is arguably very *perceptible*: it is presented for 1200 ms and subtends 0.7° of visual angle in height, and between 2° and 3.9° of visual angle in width.

Data from one experiment in our series (Aimola Davies et al., in preparation) illustrate how the full-attention-trial exclusion criterion is potentially problematic. In one condition of the experiment, the unexpected stimulus was incongruent with the observer’s attentional set, and we confidently predicted a high rate of inattentional blindness. We were, therefore, surprised to find that the rate of inattentional blindness for this condition was low, at 37.5% (6/16) – certainly lower than the rates of inattentional blindness (around 60%) that we had found in other

experimental conditions in which the unexpected stimulus was incongruent with the observer's attentional set.

This surprising finding came about because there were *seven* observers in this condition who failed to detect the unexpected stimulus on the full-attention trial. There was nothing in the primary-task performance of these seven observers on preceding trials to suggest that they were not following instructions or not engaging properly with the experiment, or that they had undiagnosed or unreported visual problems, but they were excluded in accordance with the full-attention-trial exclusion criterion. Had we not excluded observers who failed to report the unexpected stimulus on the full-attention trial, the rate of inattention blindness would instead have been 56.5% (13/23) – broadly comparable with the rate in similar experimental conditions (in which all observers had detected the unexpected stimulus on the full-attentional trial). Thus, although a majority of observers did not notice the unexpected stimulus on the critical trial, implementing the full-attention-trial exclusion criterion gave the impression that the stimulus was detected by nearly two-thirds of observers. Our decision to implement or not to implement the full-attention-trial exclusion criterion affected the results of our subsequent statistical analyses, and made a difference to the theoretical claims (about when and why inattention blindness occurs) that could be supported by those results. We also found that, when the exclusion criterion was implemented, a decision to present or not to present a forced-choice display after the critical trial made a similar difference to our results (see below, Section 4.2).

The example that we have described was not an isolated case. Several conditions in this experimental series (Aimola Davies et al., in preparation) produced high rates of inattention blindness on the full-attention trial and, in all cases, the unexpected stimulus was incongruent with the observer's attentional set. In the next sections, we investigate – empirically and conceptually – the nature and the extent of the potential problems associated with the full-attention-trial exclusion criterion.

3. Summary of 56 studies (125 experiments) investigating inattention blindness

The past decade has witnessed a steep rise in studies on inattention blindness and other failures of visual awareness (e.g., attentional blink, change blindness, motion-induced blindness, and repetition blindness). A PubMed search for “inattention blindness” identifies only 21 articles published before 2005 compared with 149 articles published between 2006 and 2016. Our aim was to review systematically how different inattention blindness studies were structured, in terms of critical, divided-attention, and full-attention trials. To be included in our

review, articles needed to present an experimental investigation of inattention blindness (as opposed to simply mentioning inattention blindness in the text), include details about the experimental trials (e.g., whether they included a critical trial, divided-attention trial, and full-attention trial), and be computer- or video-based (rather than real-world investigations). Fifty-six articles, published between 1999 and 2016, met our conditions.

Table 1 details the methods used for investigating inattention blindness. All experiments incorporated a critical trial, and the table indicates whether a forced-choice display was presented. The table also indicates whether the experiment included a divided-attention trial, whether it included a full-attention trial, and whether (and how many) observers were excluded from the data analysis for failing to report the unexpected stimulus on the full-attention trial. Here, we briefly summarise the information that we have collated in the table.

Of the 56 studies, 21 comprised one experiment; the remaining 35 studies comprised between two and eight experiments, giving a total of 125 experiments. Of the 125 experiments, 47 used a static paradigm, and 78 used a dynamic paradigm. Each experiment involved between one and 16 experimental conditions. Of the 125 experiments included in the table:

- 27 involved a critical trial, but no divided- or full-attention trials;
- 37 included a critical trial and a full-attention trial, but no divided-attention trial;
- 61 included a critical, a divided-attention, and a full-attention trial.

When probing detection of the unexpected stimulus, a forced-choice display was presented after the critical trial in 52 of 125 experiments. Of these 52 experiments:

- 6 involved a critical trial, but no divided- or full-attention trials;
- 9 included a critical trial and a full-attention trial, but no divided-attention trial;
- 37 included a critical, a divided-attention, and a full-attention trial.

There were 98 experiments that included a full-attention trial. In eight of these experiments, it was not stated whether the full-attention-trial exclusion criterion was implemented, or the number of observers who were excluded was not specified, or there was no unexpected stimulus. There were also 9 experiments in which the full-attention-trial exclusion criterion was not implemented. Thus, we focus on the remaining 81 experiments, in which there was a full-attention trial and the exclusion criterion was implemented. In 36 of these 81 experiments (16 static, 20 dynamic), all observers detected the unexpected stimulus on the full-attention trial: 17 of these 36 experiments included presentation of a forced-choice display prior to the full-attention trial (e.g., after the critical trial), and 19 did not. In the remaining 45 of the 81 experiments (15 static, 30 dynamic), between 1% and 32% of observers were excluded on the

basis of failing to detect the unexpected stimulus on the full-attention trial. Let us provide a further breakdown of these exclusions:

- 27 experiments (9 static, 18 dynamic) excluded 1-5% of observers: 13 of these experiments included prior presentation of a forced-choice display, and 14 did not;
- 9 experiments (3 static, 6 dynamic) excluded 6-10% of observers: 3 of these experiments included prior presentation of a forced-choice display, and 6 did not;
- 9 experiments (3 static, 6 dynamic) excluded more than 10% of observers: 2 of these experiments included prior presentation of a forced-choice display, and 7 did not.

To summarise, Table 1 provides a detailed overview of the experimental methods used in 56 studies of inattention blindness. The reader will appreciate the considerable variation across different research projects, in terms of the structure of experimental trials, whether a forced-choice display was presented, and whether the full-attention trial was used to screen and exclude observers. This variability is not inconsequential. In the next section, we present hypothetical experiments that demonstrate how different methodological decisions could affect the results of statistical analyses, as well as the conclusions that researchers draw.

4. *Potential problems with the full-attention-trial exclusion criterion*

4.1. *Comparing studies with and without a full-attention trial: Two hypothetical experiments*

In Table 1, there are 27 experiments (out of 125) that comprised only a critical trial, and had no divided- or full-attention trial. Studies that do not include a full-attention trial cannot, of course, exclude observers on the grounds of failing to detect the unexpected stimulus under conditions of full attention. Comparing different experimental set-ups highlights the potential impact of the researchers' decision to include or not include a full-attention trial, and to implement or not implement the full-attention-trial exclusion criterion.

Our first two hypothetical experiments – each with 20 observers – are *identical* up to the critical trial. Experiment 1 does not include a full-attention trial and all observers are included in the data analysis. Experiment 2 does include a full-attention trial and observers who fail to report the unexpected stimulus on this trial are excluded from the data analysis. In Experiment 1, 7 observers fail to report the unexpected stimulus on the critical trial. Thus the *reported* rate of inattention blindness is 35% (7 of 20 observers). In Experiment 2, as in Experiment 1, 7 observers fail to report the unexpected stimulus on the critical trial. However, all 7 observers also fail to report the unexpected stimulus on the full-attention trial. They are excluded from the data analysis and the reported rate of inattention blindness on the critical trial is *adjusted* – from 35%

(7 of 20 observers) to 0% (0 of 13 observers). When comparing Experiments 1 and 2, the *unadjusted* rates of inattentional blindness on the critical trial are not numerically different and, of course, not significantly different (35% versus 35%, $p = 1.000$), but the adjusted rates (based on implementing the full-attention-trial exclusion criterion in Experiment 2) are significantly different (35% versus 0%, $\chi^2(1, N = 33) = 3.871, p = .0491$). Thus, the reported results suggest that there was better detection of the unexpected stimulus on the *critical trial* in Experiment 2 than in Experiment 1. This interpretation of the findings is not warranted because, on the critical trials of both Experiment 1 and Experiment 2, the same proportion of observers detected the unexpected stimulus (13 of 20). We refer the reader to Table 2, which provides a summary of each of the four hypothetical experiments that we discuss.

Table 2. *Summary of our hypothetical Experiments (1-4), including information about sample size (N), unadjusted rate of inattentional blindness, full-attention-trial exclusions, and adjusted rate of inattentional blindness.*

Hypothetical Experiment	N	Unadjusted Rate of Inattentional Blindness (Critical Trial)	Full Attention Trial Exclusions	Adjusted Rate of Inattentional Blindness (Critical Trial)
1	20	7/20 – 35%	No full-attention trial	7/20 – 35%
2	20	7/20 – 35%	7	0/13 – 0%
3	20	7/20 – 35%	0	7/20 – 35%
4 (C1)	20	3/20 – 15%	0	3/20 – 15%
4 (C2)	20	11/20 – 55%	5	6/15 – 40%
4 (C3)	20	16/20 – 80%	0	16/20 – 80%

4.2. *Forced-choice displays, and detection of the unexpected stimulus on the full-attention trial*

We have shown that two hypothetical experiments with identical rates of inattentional blindness on the critical trial could lead to different conclusions, depending on whether or not they include a full-attention trial. This is also the case for two studies that both include a full-attention trial and both implement the exclusion criterion, but differ as to whether or not they present a forced-choice display to observers who do not report the unexpected stimulus on the critical trial. The next hypothetical experiment illustrates this point.

Experiment 3 is identical to Experiments 1 and 2 up to the critical trial. Once again, 7 of 20 observers fail to report the unexpected stimulus on the critical trial but in Experiment 3, a forced-choice display is presented after the critical trial and all observers report the unexpected stimulus on the full-attention trial. Consequently, despite the exclusion criterion, all observers are included in the data analysis and the reported rate of inattentional blindness is 35%, just as it was

in Experiment 1. This is significantly higher than the reported rate of inattention blindness (0%) in Experiment 2 ($\chi^2(1, N = 33) = 3.871, p = .0491$). As before, interpretation of the results as suggesting better detection of the unexpected stimulus on the critical trial in Experiment 2 than in Experiment 3 is not warranted. On the critical trials of both Experiment 2 and Experiment 3, the same proportion of observers detected the unexpected stimulus (13 of 20).

Given the number of studies that present a forced-choice display after the critical trial, this is not a purely hypothetical example. In Table 1, there are 81 (out of 125) experiments in which there was a full-attention trial and the exclusion criterion was implemented. In 35 of these 81 experiments, a forced-choice display was presented. Interestingly, in only 5 (14%) of these 35 experiments were more than 5% of observers excluded in accordance with the full-attention-trial exclusion criterion. In the remaining 46 experiments, no forced-choice display was presented and, in 13 (28%) of these experiments, more than 5% of observers were excluded.¹

Why might the inclusion of a forced-choice display help observers to detect the unexpected stimulus on the full-attention trial? Observers who have been presented, in a forced-choice display, with several candidates for the unexpected stimulus on the critical trial (e.g., a capital E, a heart, a triangle, and a cross; Most et al., 2001) may set their attention generally, for letters or shapes, or more specifically, for those particular shapes, when they come to the full-attention trial. An unexpected stimulus that matches the observer's attentional set is more likely to be detected (e.g., Most et al., 2001; 2005). And even if such an attentional set is not adopted intentionally, detection of the unexpected stimulus may be primed by its recent exposure in the forced-choice procedure. Observers who have *not* viewed a forced-choice display after the critical trial do not have the advantage of this additional information to guide their attention on the full-attention trial. Despite the experimenter's probe questions following the critical trial, and the new task instructions before the full-attention trial, these observers may, for example, continue to focus on the primary-task stimuli, perhaps looking for some subtle difference from the earlier trials.

In summary, by excluding observers who fail to report the unexpected stimulus on the full-attention trial, experimenters risk underestimating rates of inattention blindness on the critical trial in studies that do not include a forced-choice procedure (compared with studies that do include a forced-choice procedure).

¹ In our own experiments (Aimola Davies et al., in preparation), observers who fail to detect the unexpected stimulus on the critical trial, and are then presented with a forced-choice display, *never* fail to detect the unexpected stimulus on the full-attention trial. In contrast, when tested with an identical paradigm, but no forced-choice display, some observers do fail to detect the unexpected stimulus on the full-attention trial.

4.3. Comparing different conditions within a single experiment

So far we have argued that the full-attention-trial exclusion criterion may lead to problems when we compare results from different experiments – when one experiment does, but the other does not, include a full-attention trial; or one does, but the other does not, include presentation of a forced-choice display following the critical trial. Now we will argue that the exclusion criterion may *also* lead to problems when conditions are compared within a single experiment, even if the researchers have followed the same protocol across all of the experimental conditions. It is theoretically possible that application of the full-attention-trial exclusion criterion may change the outcome of statistical analyses when conditions are compared. A significant difference may be turned into a null result or *vice versa*. We illustrate the two theoretical possibilities with our final hypothetical experiment.

Experiment 4 has three experimental conditions, each with 20 observers. Although we present this as a hypothetical experiment, it was precisely the pattern of data in Conditions 1 and 2 that motivated our investigation of possible problems with the full-attention-trial exclusion criterion. In Condition 1, 3 observers fail to report the unexpected stimulus on the critical trial, but all observers report the unexpected stimulus on the full-attention trial. Thus the reported (unadjusted) rate of inattention blindness is 15% (3 of 20 observers). In Condition 2, 11 observers fail to report the unexpected stimulus on the critical trial, and the unadjusted rate of inattention blindness is 55% (11 of 20 observers). However, 5 of the observers who fail to report the unexpected stimulus on the critical trial also fail to report it on the full-attention trial, and are therefore excluded from the data analysis. The reported rate of inattention blindness on the critical trial is *adjusted* – from 55% (11 of 20 observers) to 40% (6 of 15 observers). Finally, in Condition 3, 16 observers fail to report the unexpected stimulus on the critical trial, but all observers report the unexpected stimulus on the full-attention trial. Thus the reported (unadjusted) rate of inattention blindness is 80% (16 of 20 observers).

In Experiment 4, the unadjusted rates of inattention blindness on the critical trial for Conditions 1 and 2 are significantly different (15% versus 55%, $\chi^2(1, N = 40) = 5.385, p = .0203$), but the adjusted rates are not significantly different (15% versus 40%, $p = .1992$). Thus, the full-attention-trial exclusion criterion turns a significant difference into a null result. In contrast, the unadjusted rates of inattention blindness on the critical trial for Conditions 2 and 3 are *not* significantly different (55% versus 80%, $p = .1769$), but the adjusted rates *are* significantly different (40% versus 80%, $\chi^2(1, N = 35) = 4.286, p = .0384$). Thus, the full-attention-trial exclusion criterion turns a null result into a significant difference.

4.4. *Are our hypothetical experiments too extreme?*

In two of our hypothetical experiments, observers were excluded on the basis of the full-attention-trial exclusion criterion. The rates of exclusion are: 35% (Experiment 2) and 25% (Experiment 4, Condition 2). Exclusion rates of at least 25% were reported in only two of the 125 experiments in Table 1 (Furley et al., 2010; Kreitz et al., 2015c). Indeed, of the 81 experiments that (1) included a full-attention trial, (2) implemented the full-attention-trial exclusion criterion and (3) reported the number of exclusions, only 18 experiments (22%) excluded more than 5% of observers and only 9 experiments (11%) excluded more than 10% of observers (see above, Section 3). At this point, the reader may query the relevance of our apparently extreme examples, pointing out that rates of exclusion in the literature are generally quite modest.

We have three responses to this query. First, hypothetical Experiments 1, 2, and 3 provide clear illustrations of methodological issues related to the consequences, for comparisons between experiments, of the inclusion (or not) of a full-attention trial and of the inclusion (or not) of a forced-choice display following the critical trial. Second, Conditions 1 and 2 of hypothetical Experiment 4 are conceptually analogous to conditions in our series of experiments (Aimola Davies et al., in preparation). Specifically, Condition 2 is analogous to an experimental condition for which data was presented earlier (Section 2.4). In that condition, seven of 23 observers (30%) failed to detect the unexpected stimulus on the full-attention trial. Third, we note that researchers generally report a single value indicating the total number of excluded observers across all conditions of an experiment, rather than reporting a breakdown of the number of excluded observers in each condition of the experiment. For example, it would be typical for a researcher to report that five observers were excluded from an experiment with 100 observers (5%), without reporting the breakdown of the number of excluded observers in each of perhaps five conditions ($n = 20$) of this experiment. If the five excluded observers were in the same experimental condition, the rate of exclusion for that condition could be 25%. We believe that this clustering of excluded observers is potentially problematic, but that it is also potentially informative. It may indicate a strong and pervasive form of inattentional blindness, in which an attentional set persists beyond the critical and divided-attention trials. In Table 1, we have reported the percentage of observers excluded from experiments, and we have also provided information about the number of conditions (1-16) in each experiment. What at first might seem like a small overall percentage of excluded observers, because it is spread across an entire experiment, might in fact become substantial *if* (hypothetically) the excluded observers were to cluster in a small number of conditions (e.g., in just one condition) of the experiment.

The penultimate column of Table 1 highlights this potential effect of clustering, by presenting the total number of excluded observers in each *experiment* as a proportion of observers in a single *condition* of the experiment. This illustrative hypothetical clustering of excluded observers in a single condition results in 20 of 81 experiments (25%) excluding more than 10% of observers from a condition, and indeed, 12 experiments (15%) excluding more than 25% of observers from a condition. Thus, while the hypothetical experiments that we have presented in this section may seem extreme, these rates of exclusion could quite easily occur, as a result of clustering. They may even occur in experiments with only a small overall percentage of excluded observers. For example, there are five experiments in Table 1, where the overall exclusion rate was between 1% and 5% of observers, but hypothetical clustering of exclusions in a single condition would increase the exclusion rate to over 10% (11% – 30%).

5. *Inattentional blindness on the full-attention trial*

We have suggested that there are cases of failure to notice the unexpected stimulus on the full-attention trial, clustering in experimental conditions that produce high levels of inattentional blindness on the critical trial. These cases are to be explained (we suggest), not by the observer's failure to follow instructions or failure to engage properly with the experiment, nor by an undiagnosed or unreported visual problem, but by something theoretically more interesting – a kind of inattentional blindness on the so-called 'full-attention' trial. How might this continued failure to notice the unexpected stimulus arise? A natural working hypothesis would be that factors contributing to inattentional blindness on the full-attention trial include (i) the persistence, to some degree, of the attentional set with which the unexpected stimulus is incongruent and (ii) the observer's engagement in a difficult (attention-engaging) perceptual task.

Attentional set is a setting of priorities. Visual stimuli that are congruent with the attentional set are prioritised for selection, but an observer can adopt an attentional set while looking at a blank screen and engaging in no perceptual task. Immediately before the full-attention trial, the observer is told that he or she no longer needs to perform the primary task. But the attentional set that has been adopted for the preceding trials may persist, even though the observer does not carry out the primary task. It is a question for future research which features of experimental paradigms make this persistence of attentional set more likely.

A persisting attentional set is not sufficient to account for inattentional blindness on the full-attention trial; observers' attention would also need to be engaged by a difficult perceptual task. Inattentional blindness on the critical trial is increased when the perceptual load or difficulty

of the primary task is greater (Cartwright-Finch & Lavie, 2007; Simons & Chabris, 1999). Koivisto and Revonsuo (2009) manipulated perceptual load and found that detection of an incongruent unexpected stimulus was significantly reduced when the perceptual load was higher, while detection of a congruent unexpected stimulus was not significantly affected by perceptual load. Our suggestion is that, even on the full-attention trial, if the original attentional set persists, then an observer who is performing a difficult perceptual task may fail to notice a stimulus that is incongruent with the attentional set. But what would this difficult perceptual task on the full-attention trial be?

Immediately before the full-attention trial, observers may be instructed simply to “watch the display” (Most et al., 2001, p. 11) or, more specifically, to look for “anything extra” (Cartwright-Finch & Lavie, 2007, p. 327) or “anything new or additional” (Koivisto & Revonsuo, 2007, p. 845). Some or all of these instructions may, in effect, replace one attention-engaging task with another; namely, looking for something new or additional. The term “full-attention trial” may be – at least with respect to some observers – something of a misnomer. Observers may adopt various strategies when looking for something new or additional and, despite the new instruction that there is no need to perform the primary *task*, some observers may focus on the primary-task *stimuli*, perhaps looking for some subtle difference from the earlier trials. For example, Koivisto and Revonsuo (2009) reported that observers who failed to detect the unexpected stimulus, even on the full-attention trial, often reported some (unreal) difference related to the primary task stimuli.

Why, according to this two-factor account of inattention blindness on the full-attention trial (persisting attentional set plus difficult perceptual task), does it tend to *cluster* in particular experimental conditions? Our provisional answer is that inattention blindness on the full-attention trial tends to cluster in conditions with the highest rates of inattention blindness on the critical trial because it is rather unlikely that an observer would *notice* the unexpected stimulus on the critical trial and then *fail to notice* it on the full-attention trial. Thus, the rate of inattention blindness on the critical trial provides an approximate upper bound for the rate on the full-attention trial. In conditions where the rate on the critical trial is low, the upper bound for the rate on the full-attention trial is correspondingly low.

6. Conclusion

Between the pioneering research of Mack and Rock (1998) and most subsequent work on inattention blindness, there was a shift in the rationale for, and the role of, the full-attention

trial. For Mack and Rock, the rationale for the full-attention trial was that inattentional blindness had to be distinguished from the less interesting phenomenon of blindness toward the imperceptible, and the role of the full-attention trial was to check on the perceptibility of the unexpected stimulus. In subsequent work, the perceptibility of the stimulus when attention is available has usually been assumed and the rationale for retaining the full-attention trial has been that it can be used to check that individual observers (a) understand and follow instructions, so that they engage properly with the experiment and (b) do not have undiagnosed or unreported visual problems. The role of the full-attention trial has then been to exclude from the data analysis any observers who do not meet these two conditions.

In this paper we have reviewed in some detail the use of the full-attention-trial exclusion criterion (Section 3 and Table 1) and we have pointed to some potentially problematic consequences of this use. The exclusion criterion may lead to problems when we compare results from different experiments – if one experiment does, but the other does not, include a full-attention trial (Section 4.1), or one does, but the other does not, present a forced-choice display following the critical trial (Section 4.2). It may also lead to problems when conditions are compared within a single experiment – turning a significant difference into a null result or *vice versa* (Section 4.3). Finally, we have suggested that some cases of failure to notice the unexpected stimulus on the full-attention trial should be understood as continuing inattentional blindness. We have proposed that this may arise from two factors, a persisting attentional set and performance of a difficult perceptual task that is different from the primary task (Section 5).

6.1. *The full-attention-trial exclusion criterion*

Observers who do not follow instructions and do not engage properly with an experiment should be excluded from the data analysis, as should observers with uncorrected visual problems. But observers with normal (or corrected to normal) vision, who follow instructions and engage properly with the experiment, and who show inattentional blindness on the critical trial and, for the reasons that we have indicated, also on the full-attention trial need not – and, we suggest, should not – be excluded. Their performance on the full-attention trial may contribute to our understanding of inattentional blindness itself.

If the full-attention-trial exclusion criterion is not implemented then experimenters will need to use alternative assessments to check that individual observers (a) understand and follow instructions, so that they engage properly with the experiment and (b) do not have undiagnosed or unreported visual problems. A crucial aspect of engaging with an inattentional blindness experiment is engaging with the primary task; without that engagement, inattentional blindness

would be reduced. So an assessment of primary-task performance would be relevant to condition (a) – arguably, more relevant than observers’ noticing performance on the full-attention trial, which provides no information about engagement with the primary task on the preceding trials. It would be important, though, to conceive this assessment of primary-task performance as a check, not on perfect accuracy, but on observers’ conscientiousness and commitment to *trying* to do the primary task well (Simons & Jensen, 2009). An assessment of primary-task performance would be relevant to condition (b) as well. But experimenters could also use a routine test of visual acuity. With direct assessments in place, observers could be assumed to meet conditions (a) and (b), but the full-attention trial could be retained with – not for the first time – a new rationale and role.

It is our hope that the future will see a shift in the way that the full-attention trial is utilised and reported. First, we suggest a return to full reporting of the rates of inattention blindness on the critical trial, the divided-attention trial, and the full-attention trial, in each experimental condition (as in Mack & Rock, 1998). Researchers may then choose to analyse the data from all three trials, giving particular consideration to the conditions that produce high rates of inattention blindness, even on the full-attention trial. Second, if the full-attention-trial exclusion criterion is still implemented, we suggest that experimenters adopt the practice of reporting exclusions *condition by condition*. These two suggestions are unified by the objective of maximising the information available to other researchers, and thereby enhancing our collective understanding of the factors that contribute to inattention blindness.

Table 1. *A review of 56 studies (125 experiments) investigating inattention blindness, indicating type of paradigm (static or dynamic) used in the primary-task, number of observers recruited to the experiment, number of experimental conditions, and whether the study included: a forced-choice procedure to assess implicit awareness of the unexpected stimulus (Y), a divided-attention trial (Y) and/or a full-attention trial (Y). The ninth and tenth columns provide information about the actual number (N) and percentage (%) of observers excluded from each experiment because of failure to detect the unexpected stimulus on the full-attention trial and an indicative maximum percentage of observers excluded from a condition if (hypothetically) all the observers excluded from the experiment clustered in a single condition. The final column provides additional comments.*

Reference	Paradigm	Experiment	Observers N*	Number of Conditions	Forced Choice	DA Trial	FA Trial	Observers excluded N (%)	Hypothetical clustering N (%)**	Comments
Simons & Chabris (1999)	Dynamic	1	192	16						
Most et al. (2001)	Dynamic	1	131	8	Y	Y	Y	3 (2%)	3/19 (16%)	
		2	67	4	Y	Y	Y	3(4%)	3/19 (16%)	
		3	32	2	Y	Y	Y	All detected		
Downing et al. (2004)	Static	1	144	8	Y	Y	Y	All detected		
		2	40	2	Y	Y	Y	Not specified		
Koivisto et al. (2004)	Static	1	64	4	Y	Y	Y	0		Some observers failed to detect the unexpected stimulus – none were eliminated.
		2a	35	2	Y	Y	Y	0		
		2b	35	2	Y	Y	Y	0		
Most et al. (2005)	Dynamic	1	78	4		Y	Y	13 (17%)	13/29 (45%)	
		2	30	2		Y	Y	2 (7%)	2/16 (13%)	
		3	104	4		Y	Y	2 (3%)	2/27 (7%)	
		4	44	2		Y	Y	1 (2%)	1/22 (5%)	
		5	44	2	Y	Y	Y	All detected		
		6	45	2	Y	Y	Y	All detected		
		7	43	2	Y	Y	Y	All detected		
		8	30	1		Y	Y	All detected		The unexpected stimulus was only presented on the FA trial.
Memmert (2006)	Dynamic	1	20	1						
		2	101	6						
Cartwright-Finch & Lavie (2007)	Static	1	41	2			Y	1 (2%)	1/21 (5%)	
		2	37	2			Y	1 (3%)	1/19 (5%)	
		3	40	2	Y		Y	8 (20%)	8/28 (29%)	
		4	38	2	Y		Y	2 (5%)	2/20 (10%)	

Koivisto & Revonsuo (2007)	Static	1	163	8			Y	11 (7%)	11/30 (37%)*	Total number of exclusions as a proportion of the smallest condition before exclusions = 11/17 (65%).
Memmert & Furley (2007)	Dynamic	1	31	1			Y	2 (6%)	2/33 (6%)	
		2	29	2			Y	5 (17%)	5/17 (29%)	
		3	16	1			Y	All detected		
Bressan & Pizzighello (2008)	Dynamic	1	59	2	Y	Y	Y	All detected		Forced-choice shown to noticers only.
		2	37	1	Y					Two critical trials, but observers were questioned only after the second.
		3	69	1	Y					
		4	38	1	Y					
		5	39	1	Y					
Koivisto & Revonsuo (2008)	Dynamic	1	35	4			Y	All detected		
		2	68	2			Y	3 (4%)	3/34 (9%)	
		3	32	1			Y	1 (3%)	1/32 (3%)	
		4	76	3			Y	2 (3%)	2/24 (8%)	
		5	16	1			Y	1 (6%)	1/16 (6%)	
Lee & Telch (2008)	Static	1	276	12	Y	Y	Y	0		Some observers failed to detect the unexpected stimulus – none were eliminated.
	Dynamic	2	261	12	Y	Y	Y	0		
Lo & Yeh (2008)	Static	1	90	2	Y	Y	Y	Not specified		
		2	55	2	Y	Y	Y	All detected		
Pizzighello & Bressan (2008)	Dynamic	1	90	3	Y	Y	Y	All detected		
		2	30	1	Y	Y	Y	All detected		
Railo et al. (2008)	Static	1	72	6			Y	Y	0	The order of the divided- and full-attention trials varied across observers. Some observers failed to enumerate the unexpected stimulus – none were eliminated.
		2	40	4			Y	Y	0	
White & Aimola Davies (2008)	Static	1	40	2		Y	Y	All detected		
		2	40	2		Y	Y	All detected		
		3	80	4		Y	Y	All detected		
		4	20	1		Y	Y	All detected		
		5	20	1		Y	Y	All detected		
		6	20	1		Y	Y	All detected		
Devue et al. (2009)	Static	1	48	2	Y	Y	Y	0		Some observers failed to detect the unexpected stimulus – none were eliminated.

		2	48	2	Y	Y	Y	All detected		
		3	123	3	Y	Y	Y	0		Some observers failed to detect the unexpected stimulus – none were eliminated.
Koivisto & Revonsuo (2009)	Static	1	149	8		14	Y	24 (16%)	24/40 (60%)*	Detailed information about numbers of observers in the eight conditions was not provided.
Simons & Jensen (2009)	Dynamic	1	43	1						Two blocks of trials, with one unexpected stimulus in each block.
		2	82	1						
Beanland & Pammer (2010)	Dynamic	1a	73	2	Y		Y	1 (1%)	1/37 (3%)	Two critical trials, but observers were questioned only after the second.
		1b	10	1	Y		Y			There was no unexpected stimulus in Experiment 1b.
		2	51	2	Y		Y	1 (2%)	1/26 (4%)	Two critical trials, but observers were questioned only after the second.
Furley et al. (2010)	Dynamic	1	20	1			Y	4 (20%)	4/20 (20%)	
		2	31	2			Y	6 (19%)	6/18 (33%)*	Total number of exclusions as a proportion of the smallest condition before exclusions = 6/15 (40%). However, Furley et al. provide a breakdown of exclusions by condition. The maximum proportion of exclusions was 4/16 (25%).
		3	18	1			Y	4 (22%)	4/18 (22%)	
		4	11	1			Y	3 (27%)	3/11 (27%)	
Hannon & Richards (2010)	Dynamic	1	77	1						
Hodgins & Adair (2010)	Dynamic	1	96	2						
Kuhn & Findlay (2010)	Dynamic	1	40	2			Y	All detected		
Richards et al. (2010)	Dynamic	1	68	1						
		2	82	3						
Thakral & Slotnick (2010)	Static	1	116	3	Y					
		2	110	4	Y	Y	Y	6 (5%)	6/32 (9%)	
		3	52	2	Y	Y	Y	Not specified		

Becker & Leinenger (2011)	Dynamic	1	238	6		Y	Y	2 (1%)	2/39 (5%)	
Beanland et al. (2011)	Dynamic	1	100	4			Y	Not specified		
		2b	10	1			Y	Not specified		
Chabris et al. (2011)	Dynamic	1	20	1						
		2	16	1						
		3	58	2						
Graham & Burke (2011)	Dynamic	1	112	4						
Beanland & Pammer (2012)	Dynamic	1	55	1			Y	1 (2%)	1/55 (2%)	
Bredemeier & Simons (2012)	Dynamic	1	134	1						
		2	207	1						
Hüttermann & Memmert (2012)	Dynamic	1	45	3			Y	Not specified	After the observer was informed of the unexpected stimulus, the video was replayed to check they detected it.	
		2	60	6			Y	Not specified		
Richards et al. (2012)	Dynamic	1	131	2			Y	All detected		
Pammer & Bink (2013)	Static	1	89	4	Y		Y	1 (1%)	1/23 (4%)	
Aimola Davies et al. (2013)	Dynamic	1	15	1		Y	Y	All detected		
		2a	15	1		Y	Y	All detected		
		2b	15	1		Y	Y	All detected		
		3a	15	1		Y	Y	All detected		
		3b	15	1		Y	Y	All detected		
Most et al. (2013)	Dynamic	1	87	4	Y	Y	Y	8 (9%)	8/26 (31%)	
		2	48	2	Y	Y	Y	1 (2%)	1/24 (4%)	
Wiemer et al. (2013)	Static	1	120	2		Y	Y	All detected		
Wright et al. (2013)	Dynamic	1	88	1						
Brailsford et al. (2014)	Static	1	49	2	Y	Y	Y	All detected		The authors note that observers would have been discarded if they had failed to report the unexpected stimulus on the full-attention trial
		2	49	2	Y	Y	Y	All detected		
		3	46	2	Y	Y	Y	All detected		
Calvillo & Jackson (2014)	Static	1	200	4						
Lin & Yeh (2014)	Static	1	341	16	Y		Y	All detected		
Remington et al. (2014)	Static	1	198	10			Y	11 (6%)	11/27 (41%)	
		2	199	10			Y	7 (4%)	7/23 (30%)	
		3	34	2			Y	2 (6%)	2/18 (11%)	
Kreitz et al. (2015a)	Dynamic	1	95	2	Y	Y	Y	All detected		
		2	91	2	Y	Y	Y	1 (1%)	1/46 (2%)	

		3	89	2	Y	Y	Y	All detected		
Kreitz et al. (2015b)	Static	1	114	2	Y	Y	Y	3 (3%)	3/58 (5%)	
		2	114	2	Y	Y	Y	4 (3%)	4/59 (7%)	
		3	112	2	Y		Y	6 (5%)	6/62 (10%)	
Kreitz et al. (2015c)	Static	1	988	1	Y		Y	319 (32%)	319/988 (32%)	1104 observers completed the study, but a technical problem meant that only 988 observers saw both videos. Of these 988 observers, 319 were excluded due to performance on the full-attention trial (Kreitz, personal communication).
Li et al. (2015)	Static	1	96	2	Y	Y	Y	All detected		Observers who reported seeing something additional viewed the forced-choice display.
O'Shea & Fieo (2015)	Dynamic	1	35	1		Y	Y	1 (3%)	1/35 (3%)	
Beanland & Chan (2016)	Dynamic	1	197	1			Y	2 (1%)	2/197 (1%)	
		2	148	2			Y	1 (1%)	1/71 (1%)	
Drew & Stohart (2016)	Dynamic	1	744	7						
Goldstein & Beck (2016)	Dynamic	1	227	6	Y	Y	Y	11 (5%)	11/47 (23%)	
		2	312	8	Y	Y	Y	24 (8%)	24/60 (40%)	
Horstmann & Ansorge (2016)	Static	1	64	4						
		2	48	3						
		3	48	3						
Horwood & Beanland (2016)	Dynamic	1	160	8		Y	Y	2 (1%)	2/19 (11%)	
	Static	2	158	2		Y	Y	All detected		
Kreitz et al. (2016a)	Dynamic	1	115	2		Y	Y	All detected		
Kreitz et al. (2016b)	Dynamic	2	93	1	Y	Y	Y	6 (6%)	6/93 (6%)	
Schofield et al. (2016)	Dynamic	1	794	6	Y					

* N value indicates the number of observers that the researchers included in the analysis *plus* the number of observers excluded due to the full-attention-trial exclusion criterion. We use this value to calculate the percentage (%) of observers excluded on the basis of the exclusion criterion. For example, if the researchers analysed data from 95 observers and excluded an additional 5 observers who failed to detect the unexpected stimulus on the full-attention trial, we report the percentage of excluded observers as 5% (5/100).

** To illustrate the potential impact of clustering of excluded observers, we have presented an indicative *maximum* proportion and percentage of observers that could hypothetically be excluded from a single condition. Thus, we have considered the number of observers excluded from the *experiment* as a proportion of the sum of that number and the number of observers whose data were analysed in the *smallest condition*. Consider an experiment with 100 observers, divided between five conditions. If the researchers excluded 5 observers and analysed data from 20 observers in each of four conditions and 15 observers in a fifth condition, we report the indicative maximum proportion of observers excluded from a single condition as $5/(15+5) = (25\%)$. If the researchers do not report the number of observers whose data were analysed in each condition, we use the mean: $5/(19+5) = (21\%)$.

*** In some cases, reported numbers of observers in each condition *before* exclusions due to the full-attention-trial exclusion criterion show that the *hypothetical* maximum proportion of observers excluded from a single condition could not be *actual* (because the denominator is greater than the number of observers in even the largest condition before exclusions). In these cases, we show in the Comments column the total number of exclusions as a proportion of the smallest condition before exclusions (if that information is available)

References

- Aimola Davies, A.M., Waterman, S., White, R.C., & Davies, M. (2013). When you fail to see what you were told to look for: Inattention blindness and task instructions. *Consciousness and Cognition*, 22, 221-230.
- Beanland, V., & Chan, E.H. (2016). The relationship between sustained inattention blindness and working memory capacity. *Attention, Perception, and Psychophysics*, 78, 808-817.
- Beanland, V., & Pammer, K. (2010). Looking without seeing or seeing without looking? Eye movements in sustained inattention blindness. *Vision Research*, 50, 977-988.
- Beanland, V., & Pammer, K. (2012). Minds on the blink: The relationship between inattention blindness and attentional blink. *Attention, Perception, and Psychophysics*, 74, 322-330.
- Beanland, V., Allen, R. A., & Pammer, K. (2011). Attending to music decreases inattention blindness. *Consciousness and Cognition*, 20, 1282-1292.
- Becker, M.W., & Leininger, M. (2011). Attentional selection is biased toward mood-congruent stimuli. *Emotion*, 11, 1248-1254.
- Brailsford, R., Catherwood, D., Tyson, P.J., & Edgar, G. (2014). Noticing spiders on the left: Evidence on attentional bias and spider fear in the inattention blindness phenomenon. *Laterality*, 19, 201-218.
- Bredemeier, K., & Simons, D.J. (2012). Working memory and Inattention blindness. *Psychonomic Bulletin and Review*, 19, 239-244.
- Bressan, P., & Pizzighello, S. (2008). The attentional cost of inattention blindness. *Cognition*, 106, 370-383.

- Calvillo, D.P., & Jackson, R.E. (2014). Animacy, perceptual load, and inattention blindness. *Psychonomic Bulletin and Review*, 21, 670-675.
- Cartwright-Finch, U., & Lavie, N. (2007). The role of perceptual load in inattention blindness. *Cognition*, 102, 321-340.
- Chabris, C.F., Weinberger, A., Fontaine, M., & Simons, D.J. (2011). You do not talk about Fight Club if you do not notice Fight Club: Inattention blindness for a simulated real-world assault. *i-Perception*, 2, 150-153.
- Devue, C., Laloyaux, C., Feyers, D., Theeuwes, J., & Brédart, S. (2009). Do pictures of faces, and which ones, capture attention in the inattention-blindness paradigm? *Perception*, 38, 552–568.
- Drew, T., & Stohart, C. (2016). Clarifying the role of target similarity, task relevance and feature-based suppression during sustained inattention blindness. *Journal of Vision*, 16, 1-9.
- Downing, P. E., Bray, D., Rogers, J., & Childs, C. (2004). Bodies capture attention when nothing is expected. *Cognition*, 93, B27-B38.
- Furley, P., Memmert, D., & Heller, C. (2010). The dark side of visual awareness in sport: Inattention blindness in a real-world basketball task. *Attention, Perception, and Psychophysics*, 72, 1327-1337.
- Goldstein, R.R., & Beck, M.R. (2016). Inattention blindness: A combination of a relational set and a feature inhibition set? *Attention, Perception, and Psychophysics*, 78, 1245-1254.
- Graham, E.R., & Burke, D.M. (2011). Aging increases inattention blindness to the gorilla in our midst. *Psychology and Aging*, 26, 162-166.

- Hannon, E. M., & Richards, A. (2010). Is inattentional blindness related to individual differences in visual working memory capacity or executive control functioning? *Perception*, *29*, 309-319.
- Hodgins, H.S., & Adair, K.C. (2010). Attentional processes in meditation. *Consciousness and Cognition*, *19*, 872-878.
- Horstmann, G., & Ansorge, U. (2016). Surprise capture and inattentional blindness. *Cognition*, *157*, 237-249.
- Horwood, S., & Beanland, V. (2016). Inattentional blindness in older adults: Effects of attentional set and to-be-ignored items. *Attention, Perception, and Psychophysics*, *78*, 818-828.
- Hüttermann, S., & Memmert, D. (2012). Moderate movement, more vision: Effects of physical exercise on inattentional blindness. *Perception*, *41*, 963-975.
- Koivisto, M., Hyönä, J., & Revonsuo, A. (2004). The effects of eye movements, spatial attention and stimulus features on inattentional blindness. *Vision Research*, *44*, 3211-3221.
- Koivisto, M., & Revonsuo, A. (2007). How meaning shapes seeing. *Psychological Science*, *18*, 845-849.
- Koivisto, M., & Revonsuo, A. (2008). The role of unattended distractors in sustained inattentional blindness. *Psychological Research*, *72*, 39-48.
- Koivisto, M., & Revonsuo, A. (2009). The effects of perceptual load on semantic processing under inattention. *Psychonomic Bulletin and Review*, *16*, 864-868.
- Kuhn, G., & Findlay, J.M. (2010). Misdirection, attention and awareness: Inattentional blindness reveals temporal relationship between eye movements and visual awareness. *Quarterly Journal of Experimental Psychology*, *63*, 136-146.

- Kreitz, C., Furley, P.A., & Memmert, D. (2015a). Inattention blindness is influenced by exposure time not motion speed. *Quarterly Journal of Experimental Psychology*, 69, 495-505.
- Kreitz, C., Furley, P.A., & Memmert, D., & Simons, D.J. (2016a). The influence of attention set, working memory capacity, and expectations on inattention blindness. *Perception*, 45, 386-399.
- Kreitz, C., Furley, P.A., Simons, D.J., & Memmert, D. (2016b). Does working memory capacity predict cross-modally induced failures of awareness? *Consciousness and Cognition*, 39, 18-27.
- Kreitz, C., Schnuerch, R., Furley, P.A., Gibbons, H., & Memmert, D. (2015b). Does semantic preactivation reduce inattention blindness. *Attention, Perception, and Psychophysics*, 77, 759-767.
- Kreitz, C., Schnuerch, R., Gibbons, H., & Memmert, D. (2015c). Some see it, some don't: Exploring the relation between inattention blindness and personality factors. *Plos One*, 10, e0128158.
- Lee, H.J., & Telch, M.J. (2008). Attentional biases in social anxiety: An investigation using the Inattention blindness paradigm. *Behaviour Research and Therapy*, 46, 819-835.
- Li, X., et al. (2015). Reward sensitivity predicts ice cream-related attentional bias assessed by inattention blindness. *Appetite*, 89, 258-264.
- Lin, S.-H., & Yeh, Y.-Y. (2014). Attentional load and the consciousness of one's own name. *Consciousness and Cognition*, 26, 197-203.
- Lo, S.Y., & Yeh, S.L. (2008). Dissociation of processing time and awareness by the Inattention blindness paradigm. *Consciousness and Cognition*, 17, 1169-1180.
- Mack, A., & Rock, I. (1998). *Inattention blindness*. Cambridge, MA: MIT Press.

- Mack, A., & Rock, I. (1999). Inattention blindness: An overview. *Psyche*, 5, <http://psyche.cs.monash.edu.au/v5/psyche-5-03-mack.html>
- Memmert, D. (2006). The effects of eye movements, age, and expertise on inattention blindness. *Consciousness and Cognition*, 15, 620-627.
- Memmert, D., & Furley, P. (2007). "I spy with my little eye!": Breadth of attention, Inattention blindness, and tactical decision making in team sports. *Journal of Sport and Exercise Psychology*, 29, 365-381.
- Most, S. B. (2013). Setting sights higher: Category-level attentional set modulates sustained inattention blindness. *Psychological Research*, 77, 139-146.
- Most, S. B., & Astur, R. S. (2007). Feature-based attentional set as a cause of traffic accidents. *Visual Cognition*, 15, 125-132.
- Most, S. B., Scholl, B. J., Clifford, E. R., & Simons, D. J. (2005). What you see is what you set: Sustained inattention blindness and the capture of awareness. *Psychological Review*, 112, 217-242.
- Most, S. B., Simons, D. J., Scholl, B. J., Jimenez, R., Clifford, E. R., & Chabris, C. F. (2001). How not to be seen: The contribution of similarity and selective ignoring to sustained inattention blindness. *Psychological Science*, 12, 9-17.
- Neisser, U., & Becklen, R. (1975). Selective looking: Attending to visually specified events. *Cognitive Psychology*, 7, 480-494.
- O'Shea, D.M., & Fieco, R.A. (2015). Individual differences in fluid intelligence predicts inattention blindness in a sample of older adults: A preliminary study. *Psychological Research*, 79, 570-578.

- Pammer, K., & Bink, C. (2013). Attentional differences in driving judgments for country and city scenes: Semantic congruency in inattention blindness. *Accident; Analysis and Prevention*, 50, 955-963.
- Pizzighello, S., & Bressan, P. (2008). Auditory attention causes visual inattention blindness. *Perception*, 37, 859-866.
- Railo, H., Koivisto, M., Revonsuo, A., & Hannula, M.M. (2008). The role of attention in subitizing. *Cognition*, 107, 82-104.
- Remington, A., Cartwright-Finch, U., & Lavie, N. (2014). I can see clearly now: The effects of age and perceptual load on inattention blindness. *Frontiers in Human Neuroscience*, 8, doi: 10.3389/fnhum.2014.00229
- Richards, A., Hannon, E. M., & Derakshan, N. (2010). Predicting and manipulating the incidence of inattention blindness. *Psychological Research*, 74, 513–523.
- Richards, A., Hannon, E.M., & Vitkovitch, M. (2012). Distracted by distractors: Eye movements in a dynamic Inattention blindness task. *Consciousness and Cognition*, 21, 170-176.
- Schofield, T.P., Creswell, J.D., Denson, T.F. (2016). Brief mindfulness induction reduces inattention blindness. *Consciousness and Cognition*, 37, 63-70.
- Simons, D. J. (2000). Attentional capture and inattention blindness. *Trends in Cognitive Sciences*, 4, 147-155.
- Simons, D. J. (2010). Monkeying around with the gorillas in our midst: familiarity with an inattention-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 inattention blindness for dynamic events. *Perception*, 28, 1059-1074.
- Simons, D. J., & Jensen, M. S. (2009). The effects of individual differences and task difficulty on inattention blindness. *Psychonomic Bulletin & Review*, 16, 398–403.
- Thakral, P.P., & Slotnick, S.D. (2010). Attentional inhibition mediates Inattention blindness. *Consciousness and Cognition*, 19, 636-643.
- Ward, E.J., & Scholl, B.J. (2015). Inattention blindness reflects limitations on perception, not memory: Evidence from repeated failures of awareness. *Psychonomic Bulletin and Review*, 22, 722-727.
- White, R. C., & Aimola Davies, A. (2008). Attention set for number: Expectation and perceptual load in inattention blindness. *Journal of Experimental Psychology: Human Perception and Performance*, 34, 1092-1107.
- Wiemer, J., Gerdes, A.B., & Pauli, P. (2013). The effects of an unexpected spider stimulus on skin conductance responses and eye movements: An Inattention blindness study. *Psychological Research*, 77, 155-166.
- Wright, T.J., Boot, W.R., & Morgan, C.S. (2013). Pupillary response predicts multiple object tracking load, error rate, and conscientiousness, but not inattention blindness. *Acta Psychologica*, 144, 6-11.