

Effects of Awareness on the Control of Attention

Taylor W. Webb, Hope H. Kean, and Michael S. A. Graziano

Abstract

■ Previous studies show that it is possible to attend to a stimulus without awareness of it. Whether attention and awareness are independent or have a specific relationship, however, remains debated. Here, we tested three aspects of visual attention with and without awareness of the visual stimulus. Metacontrast masking rendered participants either subjectively aware or not aware of the stimulus. Attention drawn to the stimulus was measured by using the stimulus as a cue in a spatial attention task. We found that attention was drawn to the stimulus regardless of whether or not people were aware of it. However, attention changed significantly in the absence of awareness in at least three ways. First, attention to a task-relevant stimulus was less stable over time. Second, inhibition of return, the automatic

suppression of attention to a task-irrelevant stimulus, was reduced. Third, attention was more driven by the luminance contrast of the stimulus. These findings add to the growing information on the behavior of attention with and without awareness. The findings are also consistent with our recently proposed account of the relationship between attention and awareness. In the attention schema theory, awareness is the internal model of attention. Just as the brain contains a body schema that models the body and helps control the body, so it contains an attention schema that helps control attention. In that theory, in the absence of awareness, the control of attention should suffer in basic ways predictable from dynamical systems theory. The present results confirm some of those predictions. ■

INTRODUCTION

It is now well established that attention and awareness are separable. People can attend to a stimulus in the absence of awareness of that stimulus (Norman, Heywood, & Kentridge, 2013; Hsieh, Colas, & Kanwisher, 2011; Kentridge, Nijboer, & Heywood, 2008; Koch & Tsuchiya, 2007; Jiang, Costello, Fang, Huang, & He, 2006; Tsushima, Sasaki, & Watanabe, 2006; Lamme, 2004; Lambert, Naikar, McLachlan, & Aitken, 1999; McCormick, 1997; Lambert, Beard, & Thompson, 1988). Yet attention and awareness are not entirely independent, given that interactions between them have been reported (Tsushima et al., 2006; Lambert et al., 1999; McCormick, 1997). The specific relationship between attention and awareness, however, has remained uncertain.

Recently, we proposed the attention schema theory, a possible account of the relationship between attention and awareness (Webb & Graziano, 2015; Graziano, 2013, 2014; Graziano & Webb, 2014; Kelly, Webb, Meier, Arcaro, & Graziano, 2014; Graziano & Kastner, 2011). In the theory, awareness is the internal model of attention or the attention schema. A basic principle of control theory is that a control system benefits from an internal model of the thing to be controlled (Camacho & Bordons Alba, 2004). For example, the brain constructs a body schema, an approximate internal model of the body, to help control movement. When misalignment or error in the inter-

nal model occurs, movement of the body is still possible but suffers characteristic deficits in control. If the internal model is missing, then the arm is less stably maintained in a task-relevant state, is less able to transition away from a no-longer-desired state, and is more easily perturbed by external forces (Scheidt, Conditt, Secco, & Mussa-Ivaldi, 2005; Graziano & Botvinick, 2002; Wolpert, Ghahramani, & Jordan, 1995). In the attention schema theory, the relationship between attention and awareness is similar to the relationship between the body and the body schema. In this theory, without awareness, attention should still be possible, but the control of attention should suffer. It should be less stably maintained in a task-relevant state, less able to transition away from a task-irrelevant state, and more easily perturbed by external influences.

The purpose of this study was twofold. First, because relatively few studies have directly compared attention in the presence and absence of awareness, we sought to collect basic information on how attention to a visual stimulus behaves when participants are aware of the stimulus and when they are not aware of it. The data may be useful in informing any hypothesis about the relationship between attention and awareness.

A second, more specific purpose was to test predictions of the attention schema theory. Three specific predictions were tested. Each prediction was made by applying the basic concept of an internal model in dynamical systems control to the case of attention. The first prediction concerns stability. In control theory, the internal model helps to stabilize the system in a desired state

(Camacho & Bordons Alba, 2004). The loss of an internal model leads to a reduction of stability. In the case of the arm, without an internal model, the system is less able to monitor and thus maintain a stable desired arm position. The arm wobbles. If the attention schema theory is correct, then without awareness of the stimulus, attention drawn to a task-relevant stimulus should become less stable in time. Attention should wobble. To test this prediction, in Experiment 1, we directly compared attention to a visual stimulus with and without awareness of the stimulus. We used a Posner spatial attention paradigm (Posner, 1980) in human participants to measure attention drawn by a brief visual cue. Attention was measured at 5 time points in the first 600 msec after cue presentation to provide a temporal profile. Metacontrast masking (Breitmeyer & Öğmen, 2006) was applied such that participants reported being subjectively aware of the cue on some trials and unaware of it on other, interleaved trials. In this manner, we tested how the time course of attention changed depending on the presence or absence of awareness of the cue. The specific prediction of the attention schema theory was that attention drawn by the cue would show greater stability through time in the presence of awareness and show significantly greater temporal instability in the absence of awareness.

The second prediction concerns the control of movement away from a task-irrelevant or nondesired state. In control theory, the loss of an internal model compromises a control system's ability to register when it is in a nondesired state and therefore impairs the system's ability to actively leave or avoid that nondesired state. In the case of the arm, without an internal model, the system cannot register that the arm has entered a task-irrelevant configuration, and therefore, compensatory movements away from that nondesired state are impaired. In the case of attention, there is a well-studied phenomenon in which attention actively transitions away from a task-irrelevant state. If attention is drawn to the onset of a task-irrelevant stimulus, that attention is typically rapidly attenuated and can even become "negative" in the sense that attention briefly avoids that location in preference for other locations, a phenomenon known as inhibition of return (IOR; Posner, Rafal, Choate, & Vaughan, 1985; Posner & Cohen, 1984). If the attention schema theory is correct, then without awareness of the visual stimulus, IOR should be reduced or eliminated. To test this prediction, we used the same paradigm as for Experiment 1, with one change. In Experiment 2, participants were not asked at the end of each trial whether they were aware of the cue. In this experiment, therefore, the cue became entirely behaviorally irrelevant. There was no longer any task-related reason for the participant's attention on the cue to be maintained. This condition of behavioral irrelevance is the necessary condition for IOR. The specific prediction of the attention schema theory is that, when the metacontrast masking was timed to allow awareness of the cue, IOR should be present. When the

metacontrast masking was timed to prevent awareness of the cue, IOR should be reduced or absent. This experiment is of interest beyond testing a specific prediction of the attention schema theory. Whether IOR occurs in the absence of awareness is a basic question about the behavior of attention, and it has been studied and debated before (Ivanoff & Klein, 2003; Lambert et al., 1999; McCormick, 1997). Previous experiments arguably did not use directly comparable stimuli on interleaved aware and unaware conditions to determine quantitatively how IOR may change. We hope therefore that this study will add to that literature.

The third prediction concerns perturbations by external influences. In control theory, without an internal model, the item being controlled is less internally stabilized and more affected by external drivers. In the case of the arm, without an internal model of the arm to help stabilize it, the control mechanism is less able to resist external forces. The arm becomes more easily perturbed. In the case of attention, the external influences are bottom-up, sensory-driven factors such as the luminance contrast of a visual stimulus. If the attention schema theory is correct, these external drivers of attention should have a greater impact in the absence of awareness than in the presence of awareness. Attention normally depends on stimulus contrast, but that dependence should have a steeper slope when awareness is absent than when awareness is present. To test this prediction, in Experiment 3, we again used a Posner paradigm to measure attention drawn to a visual cue. To focus on bottom-up attention, we tested attention 50 msec after the cue onset. The luminance contrast of the cue was varied, and the effect of contrast on attention was measured. The specific prediction of the theory was that stimulus contrast would have a greater impact on attention in the absence of awareness of the cue than in the presence of awareness of the cue. Once again, this experiment may have interest beyond testing a specific prediction of the attention schema theory. Many previous studies have demonstrated stimulus-driven attention in the absence of awareness of the stimulus (e.g., Tsushima et al., 2006; Lambert et al., 1999; McCormick, 1997). To our knowledge, however, no prior experiment has tested how awareness might affect the specific relationship between stimulus contrast and stimulus-driven attention. The experiment can therefore add to the growing knowledge of how attention behaves with and without awareness.

It is important to note that if the three predictions tested here are confirmed, they do not prove the attention schema theory. In testing a theory, one makes and tests a series of predictions. Each test is an opportunity to disconfirm the theory. If the data support the prediction, the theory survives as a possibility, but it is not proved because there may be alternative explanations for the same result. Only by testing a very large number of predictions over many publications will we begin to gain some confidence in the theory—or disconfirm it. In the

process of testing these predictions, we will hopefully gain valuable information that can inform any theory on the relationship between attention and awareness.

EXPERIMENT 1

Methods

Participants

Seventy-five participants were tested (44 women, 18–48 years old, normal or corrected-to-normal vision). All participants provided informed consent, and all procedures were approved by the Princeton institutional review board.

Behavioral Paradigm

Participants sat 30 cm from the monitor and used a chin rest to stabilize the head. Stimuli were presented with the MATLAB Psychophysics Toolbox. Figure 1 shows the behavioral paradigm. Each trial began with a white central fixation point on a black background. Participants were instructed to fixate during the trial. After 1 sec, the cue period began and lasted 4 refresh cycles (~50 msec). Throughout the cue period, the cue (a white spot 1.1° in diameter) was presented 6° to the left of fixation (1/3 of trials), to the right of fixation (1/3 of trials), or not presented (1/3 of trials). The cue was followed by a mask. In the long cue/mask interval condition (1/2 of trials), a cue/

mask interval of 4 refresh cycles (~50 msec) was inserted between the cue and the mask. Thus, the time from cue onset to mask onset was ~100 msec, intended to allow the cue to be subjectively visible to the participants. In the short cue/mask interval condition (1/2 of trials), the mask immediately followed the cue. Thus, in this trial type, the time from cue onset to mask onset was ~50 msec, intended to render the cue subjectively invisible to the participants. During the mask period, on all trials, two metacontrast masks (Breitmeyer & Ögmen, 2006) were presented, 6° to the left and right of fixation. The masks themselves therefore did not preferentially attract bottom-up attention to one side. Each mask consisted of a white ring with an outer diameter of 1.5° and an inner diameter of 1.1°. The duration of the mask period varied between 115 and 540 msec, depending on the total duration of the trial as described next.

After the mask period, a target was added to the already-present white circles that composed the mask. The target was either on the left (1/2 of trials) or right (1/2 of trials) of fixation. The location of the cue did not predict the location of the target. The target was presented at one of five times after cue onset: 12 refresh cycles (~165 msec), 20 refresh cycles (~270 msec), 28 refresh cycles (~380 msec), 36 refresh cycles (~485 msec), and 44 refresh cycles (~590 msec). Because of these varied target presentation times, the mask period between the cue and the target also varied accordingly. The target consisted of one white line segment extending from the top of the mask ring and one white line segment extending from the bottom of the mask ring. The segments were collinear, forming an implied line through the ring. The line was tilted toward the left (1/2 of trials) or toward the right (1/2 of trials) by 3°. Participants were required to discriminate the orientation by pressing the F key on a keyboard if the line was tilted to the left and the J key if it was tilted toward the right. The target period lasted 1 sec, and participants were instructed to respond as quickly as possible during that period. The short target response period encouraged a speeded response. Participants performed the discrimination task with a mean accuracy of 90% ($SD = 8\%$) and a mean latency of 673 msec ($SD = 127$ msec). The 10% of trials with incorrect responses or no responses during the 1 sec response window were excluded from analysis. Trials with a latency < 300 msec were also excluded because the RT was too short to be a plausible response to the discrimination task, and therefore, those trials probably represented mistaken key presses.

On each trial, after the 1-sec target period, the fixation point and all other stimuli disappeared, and a question was presented on the screen: “Did you see a circle? Y/N.” Participants were instructed that this question referred to the cue stimulus presented at the beginning of some trials. Participants pressed the F key to indicate that they had seen the cue on that trial and the J key to indicate that they had not. The awareness probe remained on the screen for 2 sec, and participants were instructed to

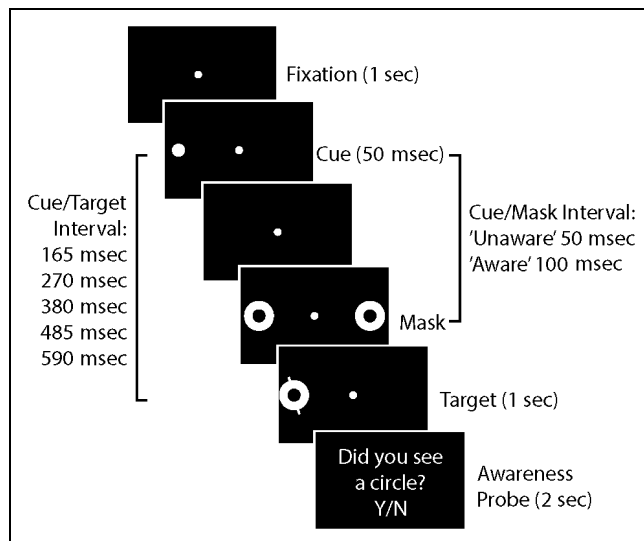


Figure 1. Paradigm for Experiment 1. Fixation was followed by cue period, premask period, mask period, target period, and awareness probe. In long cue/mask interval trials (“aware” condition), the mask was timed to allow participants to see the cue. In short cue/mask interval trials (“unaware” condition), the mask was timed to prevent participants from seeing the cue. The tilted discrimination target was presented on the same side as the cue (spatially matching as shown here) or on the opposite side as the cue (mismatching). After indicating the tilt of the target, participants were probed whether or not they had seen the cue.

respond during that time. On <1% of trials, participants failed to respond during that time window, and those trials were excluded from analysis. The mean awareness RT was 580 msec ($SD = 322$ msec). Choosing a cutoff for the minimum RT is more difficult for the awareness response. Note that the cue was presented early in the trial, and participants had as much as 1.5 sec (depending on trial type) before the awareness probe to decide if they had seen the cue. Moreover, the awareness probe was presented at a predictable time 1 sec after the target onset. Therefore, participants could, in principle, anticipate the timing of the question and respond with extremely short latency. This may account for the large spread in latency to the awareness probe. On the basis of the distribution of RTs, we chose to exclude trials with a cue awareness RT of <50 msec.

Each participant performed practice trials followed by nine runs of 24 trials each (216 trials total). Target orientation (tilted left or right), target location (left or right), cue condition (left, right, or no cue), and cue/mask interval (long or short) were randomized and counter-balanced within each run, and the cue did not predict the location of the target. The design was a mixed, within/between-subject design. The aware versus not aware manipulation (long or short cue/mask interval) was within participants. However, different participants were tested for each of the five time intervals between the cue and the target. The reason was that to test each participant at all 5 time points would have required an estimated 7.5 hr of testing per participant, which was prohibitive on practical grounds. Thus, we tested each participant with only one cue/target time interval. For each of the five cue/target time intervals, 15 participants were tested, for a total of 75 participants.

Analysis

To determine whether the division of trials into “aware” and “unaware” was justified, the participants’ responses to the awareness probe were analyzed. Responses to the awareness question from the 1/3 of trials on which no cue was presented were used to assess participants’ false positive rate, the rate at which participants indicated that a cue was present when none was. The rate was low (mean = 11%, $SD = 14\%$), indicating that participants were not guessing about the presence of the cue. Responses to the awareness question from the 2/3 of trials on which a cue was presented were used to compute cue detection rates. To assess whether the meta-contrast manipulation had the intended effect on awareness, these cue detection rates were analyzed in a 5×2 mixed factorial ANOVA: 5 Time points (between-subject variable) \times 2 Awareness conditions (within-subject variable). A highly significant main effect of Cue/mask interval confirmed that participants tended to report awareness of the cue on long cue/mask interval trials (awareness reported on 80% of trials) and tended to report no aware-

ness of the cue on short cue/mask interval trials (awareness reported on 30% of the trials; $F = 371, p = 2 \times 10^{-16}$). No other significant effects were found for the cue detection rate, showing that performance was stable across the 5 time points. Thus, not only did the manipulation separate most trials into “aware” and “unaware” but also that separation was stable across all five cue/target intervals. This confirmation allowed for the main analysis described next.

In a Posner task with high-accuracy performance, accuracy is typically an insensitive measure of attention, whereas latency is the preferred measure. The logic of the Posner task is that if the cue draws attention to itself and the subsequent target appears at the same location as the cue, then the response to the target will be faster. Likewise, if the cue draws attention to itself and the subsequent target appears at the opposite location to the cue, then the response to the target will be slower as attention is shifted from the cue to the target. To quantify this effect of the cue on attention, we used a difference of response latencies: $\Delta t = [\text{average target RT on cue/target mismatch trials}] - [\text{average target RT on cue/target match trials}]$. A larger Δt indicates that the cue drew more attention to itself. For each participant, two Δt scores were computed, one for long cue/mask interval trials (aware trials) and one for short cue/mask interval trials (unaware trials). The data were then analyzed in a 5×2 mixed factorial ANOVA: 5 Time points (between-subject variable) \times 2 Awareness conditions (within-subject variable) to determine whether the time course of attention was the same or different depending on the presence or absence of awareness.

Results

Experiment 1 tested the time course of attention to a brief visual cue over the 600 msec after cue presentation. The time course was tested with and without awareness of the cue. Figure 2A shows the results. The red line shows the results for the aware condition, and the blue line shows the results for the unaware condition. Although attention was present without awareness, it did not behave in the same way. Moreover, the change in attention was more complex than a simple increase of attention in the aware condition. The change therefore cannot be attributed to the stimulus simply having greater perceptual signal strength in the aware condition. Indeed, at 1 time point (270 msec), attention was actually greater without awareness than with it. The most striking change was that attention was less stable over time in the unaware condition than in the aware condition. The results therefore confirmed Prediction 1.

The data were analyzed with a 5×2 mixed factorial ANOVA. The main effect of Awareness was not significant ($F = 1.11, p = .30$), showing that awareness did not cause an overall increase or decrease in attention. The main effect of Time was significant ($F = 4.11, p = .005$), showing

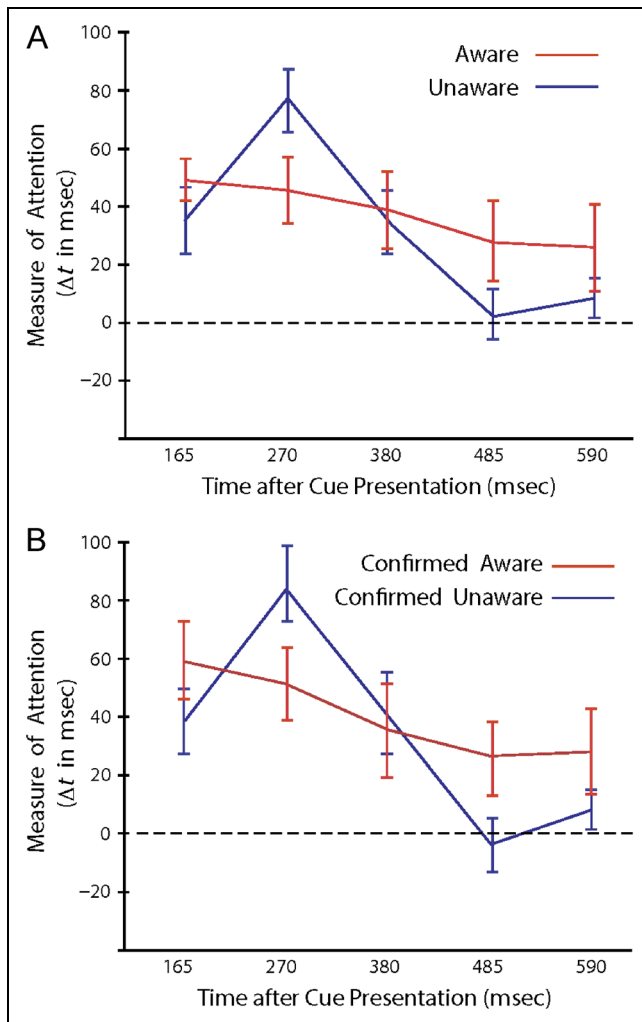


Figure 2. Results of Experiment 1. (A) Red line shows trials when the mask was timed to allow the cue to be seen. Blue line shows trials when the mask was timed to prevent the cue from being seen. The X axis shows time after cue onset. The Y axis shows attention drawn to the cue ($\Delta t = [\text{mean RT for spatially mismatching trials}] - [\text{mean RT for spatially matching trials}]$). Error bars are standard error. (B) A more selective subset of the data. Red line shows data when the mask was timed to allow the cue to be seen, and participants confirmed that they had seen it. Blue line shows data when the mask was timed to prevent the cue from being seen, and participants confirmed that they had not seen it.

that attention changed as a function of time. The interaction was also significant ($F = 3.38, p = .02$), confirming that the two curves had significantly different shapes. A 5×1 ANOVA on aware trials (Figure 2A, red line) revealed no significant effect of Time ($F = 0.67, p = .61$), indicating that attention was relatively stable during the trial in the presence of awareness. In contrast, a 5×1 ANOVA on unaware trials (Figure 2A blue line) revealed significant variation over Time ($F = 9.1, p = 5 \times 10^{-6}$). The primary difference between the aware and unaware conditions in this task, therefore, was that, in the presence of awareness, attention to the cue was more stable throughout the tested time period.

Figure 2B shows the same result as in Figure 2A, but with a selective subset of the data. Here, the “confirmed aware” trials meet two constraints: They are the 80% of trials with a long cue/mask interval on which the participant also explicitly indicated that the cue had been seen. Likewise, the “confirmed unaware” trials are the 70% of trials with a short cue/mask interval on which the participant also indicated that the cue had not been seen. This more selective data set shows a nearly identical pattern to the one in Figure 2A. The same statistical effects were found (5×2 ANOVA, no significant main effect of Awareness, $F = 0.82, p = .37$; significant main effect of Time, $F = 5.8, p = .0004$; significant interaction, $F = 2.82, p = .03$; 5×1 ANOVA, no significant effect of Time in aware trials, $F = 1.14, p = .35$; significant effect of Time in unaware trials, $F = 10.1, p = 1 \times 10^{-6}$).

EXPERIMENT 2

Methods

In Experiment 2, we tested 75 new participants not tested in Experiment 1 (56 women, 18–42 years old, normal or corrected-to-normal vision). The methods were identical to those of Experiment 1 with one exception. After the target period, participants were not asked whether they had seen the cue. The cue was therefore rendered task-irrelevant.

Results

Experiment 2 tested the effect of awareness on IOR. A task-irrelevant stimulus is needed to evoke IOR. Yet in Experiment 1, the cue was task relevant. The task relevance of the cue is not immediately obvious. The cue was not relevant to the discrimination task in a manner traditional in the Posner paradigm, because the cue did not predict the location of the target. However, the cue was still relevant to the participants’ behavior because at the end of the trial participants were asked to report whether they had seen the cue. The task therefore required participants to direct at least some attention to the cue. To modify Experiment 1 and make the cue entirely task irrelevant requires that participants not be asked about the cue at the end of the trial. Removing the end-of-trial awareness question, however, poses a challenge. The question provides a direct measure of the participant’s awareness. Can we dispense with that measure? The results of Experiment 1 show that, on most trials, the mask succeeded in rendering the cue seen on long cue/mask interval trials and unseen on short cue/mask interval trials. Moreover, using the participants’ own reports on subjective awareness to further specify the trials did not change the pattern of results (compare Figure 2A and B). Thus the trial-by-trial awareness report can be eliminated, allowing for the design of Experiment 2.

Figure 3 shows the results. In the aware condition (red line), the cue drew attention initially. At 165 msec, the

measure of attention was significantly above zero. However, unlike in Experiment 1 when the stimulus was behaviorally relevant and attention was relatively sustained, here with a task-irrelevant stimulus attention was attenuated over time. By 485 msec, attention showed evidence of active suppression with a classic IOR negative phase that recovered by 590 msec. In contrast, in the unaware condition (blue line), the negative phase of IOR was absent. The attention measure did not dip below zero. The initial attention to the cue in effect ebbed away, falling back to zero without evidence of active suppression. Prediction 2 was confirmed: IOR, the redirection of attention away from an irrelevant stimulus, was impaired without awareness.

A 5×2 mixed factorial ANOVA was performed. The main effect of Awareness was not significant ($F = 0.99, p = .32$), showing that awareness did not cause an overall increase or decrease in attention. The main effect of Time was significant ($F = 12.71, p = 8 \times 10^{-8}$), showing that attention changed as a function of time. The interaction was also significant ($F = 3.2, p = .02$), confirming that the two curves had significantly different shapes. The negative IOR in the aware condition at 485 msec was significantly below zero ($t = 2.9, p = .01$) and significantly below the unaware condition (paired t test, $t = 2.36, p = .03$).

COMPARISON OF EXPERIMENTS 1 AND 2

Figure 4 shows the data from Experiments 1 and 2 replotted for comparison. Figure 4A shows all results from aware trials. The thick line shows trials when the cue was task-relevant (participants were asked about their awareness of the cue after each trial), and the thin line shows

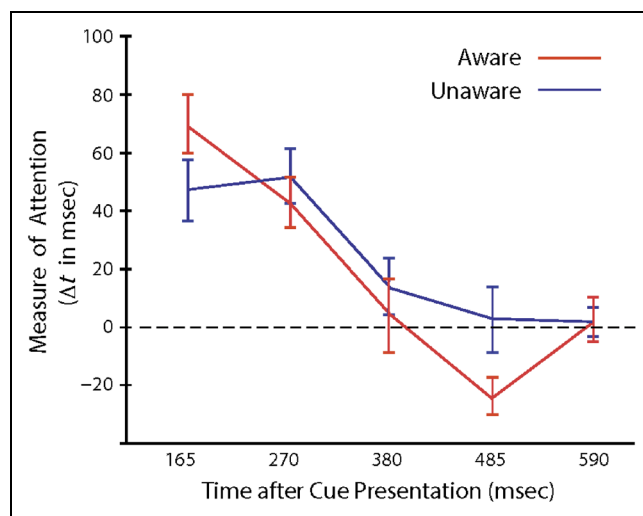


Figure 3. Results of Experiment 2. Attention to a behaviorally irrelevant stimulus. Red line shows trials when the mask was timed to allow the cue to be seen. Blue line shows trials when the mask was timed to prevent the cue from being seen. The x axis shows time after cue onset. The y axis shows attention drawn to the cue. Error bars are standard error. Only aware trials dipped below zero.

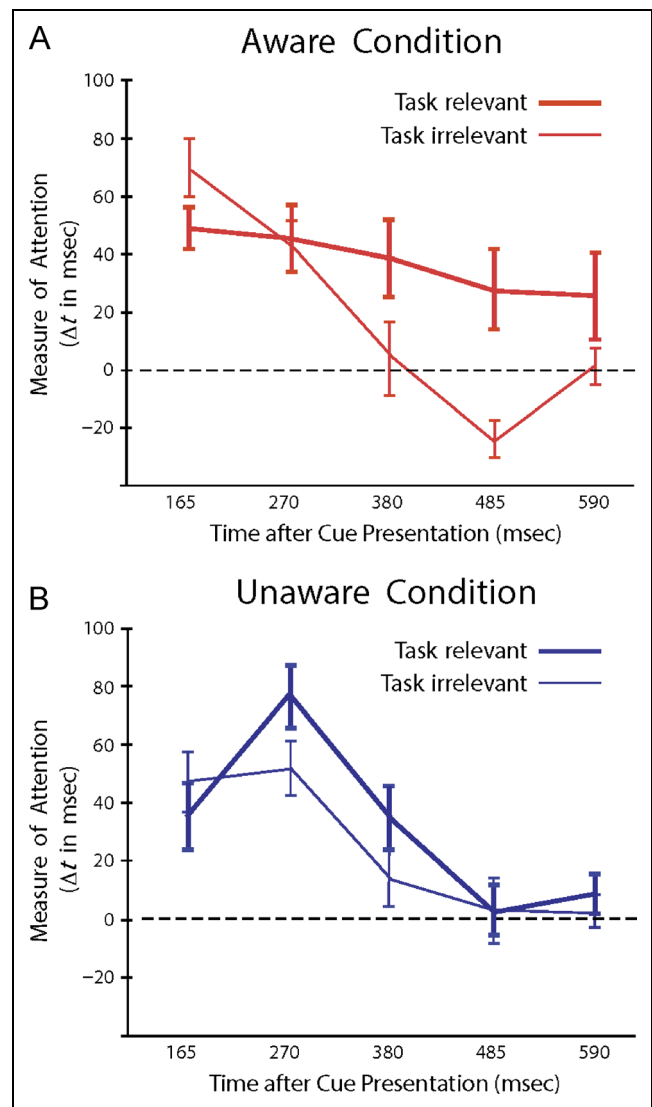


Figure 4. Comparison of Experiments 1 and 2. (A) Aware trials, when the mask was timed to allow the cue to be seen. Thick line shows trials when the cue was task relevant (participants reported at the end of each trial whether they had seen it). Thin line shows trials when the cue was task irrelevant (participants were not required to respond to it). Error bars are standard error. (B) Unaware trials, when the mask was timed to prevent the cue from being seen. Thick line shows trials when the cue was task relevant. Thin line shows trials when the cue was task irrelevant.

trials when the cue was task irrelevant (participants were not asked about the cue). These two curves differ markedly. Two characteristics stand out. First, when the cue was task relevant, attention was relatively sustained over the tested time period. Second, when the cue was task irrelevant, the attention initially drawn to the cue was attenuated and driven negatively into IOR. With awareness of the cue, attention to the cue was controlled in a task-relevant manner, whether the task required sustained attention or suppressed attention.

Figure 4B shows the results from unaware trials. The thick line shows trials when the cue was task relevant, and the thin line shows trials when the cue was task

irrelevant. The two curves are not significantly different at any time point. Most importantly, attention was not sustained when the cue was task relevant, nor was attention driven negatively into IOR when the cue was task irrelevant.

An omnibus analysis ($5 \times 2 \times 2$ mixed factorial ANOVA: 5 Time points \times 2 Awareness conditions \times 2 Task relevance conditions) revealed no significant main effect of Awareness ($F = 0.03, p = .87$), a significant main effect of Time ($F = 13.93, p = 1 \times 10^{-9}$), and a significant three-way interaction ($F = 3.27, p = .01$).

EXPERIMENT 3

Methods

We tested 26 participants not tested in the previous experiments (19 women, 18–50 years old, normal or corrected-to-normal vision).

Experiment 3 did not use the metacontrast paradigm of Experiments 1 and 2 because the mask that manipulates awareness requires time to become effective, whereas we sought to test the initial stimulus-driven attention. We therefore used a different paradigm that did not depend on masking and allowed us to test attention 50 msec after stimulus onset. Figure 5 shows the behavioral paradigm. At the start of each trial, a central fixation cross and two black boxes with their centers displaced 5° to the left and right of fixation appeared on a gray screen. Participants were instructed to fixate centrally during the trial. After 1500 msec, a 50-msec cue period began. The cue was a single pixel presented at one of two luminance levels. These two luminance levels were tailored to each

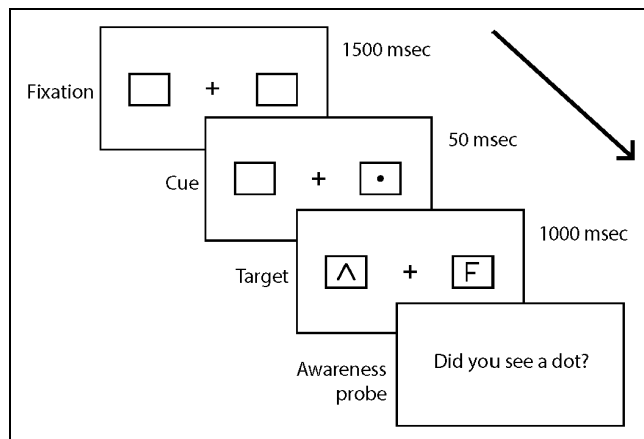


Figure 5. Behavioral paradigm for Experiment 3. Fixation onset was followed by cue period, target period, and awareness probe. Low contrast cues were presented on 2/5 of trials; high contrast cues were presented on 2/5 of trials. One fifth were catch trials with no cue. The target could be on the same side as the cue (spatially matching as shown here) or the opposite side (mismatching). Participants indicated the identity of the target (“F” or “A”) as quickly as possible while ignoring the distractor. Participants were then probed whether or not they had seen the cue on a scale of 1 (*sure it was not visible*) to 7 (*sure it was visible*).

participant based on performance in practice runs. For both luminance levels, the cue was near perceptual threshold such that on some trials participants reported seeing it and on some trials reported not seeing it. This trial-by-trial variation provided the “aware” and “unaware” trials without having to rely on a mask. With respect to the gray background (RGB of 190/190/190), the higher contrast cue was chosen from the luminance range of RGB between 50/50/50 and 110/110/110, and the lower contrast cue was chosen from the luminance range of RGB between 145/145/145 and 170/170/170. The cue was presented at low signal strength in the center of the left box on 1/5 of trials, at low signal strength in the center of the right box on 1/5 of trials, at high signal strength in the center of the left box on 1/5 of trials, and at high signal strength in the center of the right box on 1/5 of trials. On 1/5 of trials, no cue was presented.

The cue period was followed by a 1-sec target period during which two items were presented, one in each box. The items consisted of one target character, an “A” or “F,” and one distractor character, identical to an “A” or “F” but with the omission of a horizontal line. Participants were instructed to ignore the distractor and identify the target using the “A” and “F” keys on a keyboard as quickly as possible.

Following the target period, the screen displayed the question: “Did you see a dot?” Participants rated on a scale of 1–7 how certain they were that they saw a dot before the target. A 1 indicated confidence a dot was *not* seen. A 7 indicated confidence a dot *was* seen. Participants responded with either 1 or 7 on most trials (81%, 40% 1 and 41% 7). To ensure that we analyzed only trials in which participants were definitely aware or not aware of the dot, we limited analysis to trials in which participants indicated either 1 or 7.

Each participant performed 24 runs of 20 trials each (480 trials total) completed in a single session. Trial types were counterbalanced and randomized within each run. The side on which the cue appeared, the signal strength of the cue, and the side on which the target appeared were randomized and the location of the cue did not predict the location of the target.

Participants performed the cue awareness response with low false alarm rate (rate at which they reported 7 when no cue was present, mean = 2%, $SD = 2\%$) and performed the target discrimination task with high accuracy (mean = 95%, $SD = 3.4\%$). The same Δt metric as in Experiment 1 was used as a measure of attention. For each participant, four Δt scores were computed corresponding to a 2×2 design: Aware versus unaware \times Low contrast versus high contrast. The data were then analyzed with a 2×2 repeated-measures ANOVA.

Results

Experiment 3 tested stimulus-driven attention 50 msec after the onset of the stimulus. The effect of stimulus contrast on attention was measured. Figure 6 shows

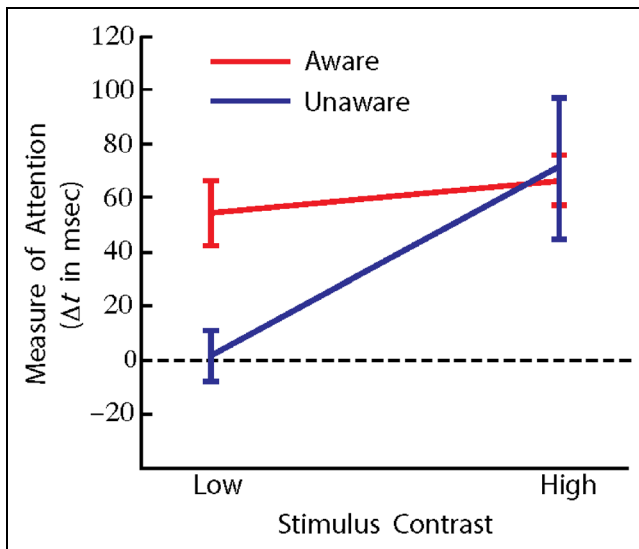


Figure 6. Results of Experiment 3. Red line shows trials when participants reported being aware of the cue. Blue line shows trials when participants reported being unaware of the cue. The x axis shows cue contrast. The y axis shows attention drawn to the cue measured by $\Delta t = [\text{mean RT for spatially mismatching trials}] - [\text{mean RT for spatially matching trials}]$. Error bars are standard error.

the results. When participants reported being definitely aware of the cue (red line), attention to the cue was weakly modulated by contrast. When participants reported being definitely unaware of the cue (blue line), attention to the cue was strongly modulated by contrast. Because the results include an interaction, they cannot easily be explained in terms of simple main effects such as an overall increase of attention in aware trials or an overall increase in signal strength in aware trials. The results confirmed Prediction 3: Stimulus-driven attention was more affected by stimulus contrast in the absence of awareness than in the presence of awareness.

A 2×2 repeated-measures ANOVA showed no significant main effect of awareness ($F = 1.2, p = .28$) indicating that Awareness was not associated with an overall increase or decrease in attention. The main effect of Stimulus contrast was significant ($F = 9.1, p = .005$), confirming that attention was overall greater to the higher contrast stimulus. Crucially, the interaction was significant ($F = 5.2, p = .03$), confirming that stimulus contrast had a greater effect on attention in the unaware condition than in the aware condition.

GENERAL DISCUSSION

This set of experiments tested several aspects of visual attention in the presence and in the absence of awareness of the visual stimulus. The main finding is that, although attention is possible without awareness, it changes in complex ways. Because these experiments directly compared aware and unaware conditions in inter-

leaved trials, it was possible to measure these changes in quantitative detail. A complex relationship between attention and awareness has been demonstrated before (e.g., Hsieh et al., 2011; Tsushima et al., 2006; Lambert et al., 1999; McCormick, 1997). The present report adds to the specific information about how attention changes when awareness is removed.

Experiment 1 plotted the time course of attention to a behaviorally relevant stimulus, during the 600 msec after stimulus presentation. Although attention was possible without awareness, the time course changed. Attention became more variable over time in the absence of awareness. Because of this variability, at 1 time point attention was actually greater when the participant was unaware of the stimulus. Others have also reported a greater attention to an unaware stimulus (e.g., Tsushima et al., 2006; McCormick, 1997). The present findings suggest that attention is not generally increased or decreased in the absence of awareness but rather becomes less stable in time and therefore, at some time points, may rise above or fall below the level it would have had in the presence of awareness.

In Experiment 2, the time course of attention to a behaviorally irrelevant stimulus was plotted. Again, although attention was possible without awareness, the time course differed from in the aware condition. When participants were aware of the stimulus, they demonstrated IOR, actively suppressing attention to the stimulus at the 485-msec time point. When participants were unaware of the stimulus, IOR was no longer observed. Whether IOR to a stimulus is present without awareness of the stimulus has been debated (Ivanoff & Klein, 2003; Lambert et al., 1999; McCormick, 1997). The present results do not necessarily resolve that debate. Perhaps some IOR is still present in some conditions or at a magnitude too small for our experiment to detect. By directly comparing the aware and unaware conditions, however, the present results do suggest that IOR is at least reduced in the absence of awareness.

In Experiment 3, stimulus-driven attention 50 msec after the onset of the visual stimulus was measured while stimulus contrast was varied. Stimulus contrast had a greater effect on attention when participants were not aware of the stimulus than when they were aware of it. This finding adds to the growing literature showing that not just top-down attention but also stimulus-driven attention is affected by the presence or absence of awareness (Lambert et al., 1999). In our experiment, because attention was tested within 50 msec of cue onset, there was presumably no time for any substantial cognitive or top-down involvement in regulating how much attention was drawn to the stimulus. Yet this fast, bottom-up attention was still affected by the presence or absence of awareness.

The three experiments presented here were motivated by the attention schema theory, one possible account of the relationship between attention and awareness. The

theory is based on concepts from dynamical systems control. From it, we made three specific predictions about the behavior of attention with and without awareness. All three predictions were confirmed. For a task-relevant stimulus, attention became less stable when awareness was absent (Experiment 1). For a task-irrelevant stimulus, IOR was reduced when awareness was absent (Experiment 2). Stimulus-driven attention was more affected by stimulus contrast when awareness was absent (Experiment 3). These three tests of course do not prove the theory. Rather, the theory would have been disconfirmed had the predictions failed. The attention schema theory remains one possible account of the complex relationship between attention and awareness.

In a possible alternative interpretation, one might say, “If I’m aware of the stimulus, then of course I can consciously decide to pay more or less attention to it. In that trivial way, awareness contributes to the control of attention.” We term this the cognitive control hypothesis. It is distinct from the hypothesis proposed here, the attention schema hypothesis. The cognitive control hypothesis is tempting because it is intuitive, but it is an unlikely explanation. First, it fails to explain. It asserts with some circularity that consciousness of a stimulus comes with the capability for conscious control. Second, although it might apply to the results of Experiment 1, it is unlikely to account for Experiments 2 and 3. In Experiment 2, it is unlikely that participants made a conscious cognitive decision to employ IOR. IOR is usually considered to be an automatic process, not the result of a high-level cognitive decision. In Experiment 3, the stimulus-driven attention effects were obtained within 50 msec of stimulus onset, presumably too fast to be the result of a cognitive decision. The results are more consistent with the attention schema theory. Without awareness of the stimulus, many aspects of attentional control changed including fast automatic aspects of control. The changes in control were consistent with the loss of an internal control model of attention. Of course other, more successful alternative explanations may emerge as more experiments are conducted.

The attention schema theory may shed light on a recent line of experiments in which the stability of a visual representation in the brain is greater when the participant is aware of the stimulus (Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene, 2015; Schurger, Pereira, Treisman, & Cohen, 2010). In those studies, the neural signature of visual awareness is a cortical pattern of activity that is more stable across time within a trial and more consistent across trials. In the attention schema theory, an internal model would help to stabilize attention and therefore stabilize the visual representation. However, it is also important to note that in the attention schema theory, stability is only one aspect of the performance of the system. Awareness should enhance stability when stability is task relevant and should enhance the agility of the system when change is useful.

A central question in consciousness research is whether subjective awareness serves any function or is an epiphenomenon. In the attention schema theory, subjective awareness serves a specific mechanistic function. It is the internal model for the dynamical systems controller that guides attention. Without it, attention is still possible and even some aspects of attentional control remain, but the brain is no longer capable of a nuanced, model-based control. Because attention is one of the most consequential processes in the brain—it determines what information is deeply processed and therefore which signals affect behavior and memory—the attention schema theory suggests that awareness is no epiphenomenon but lies at the center of our ability to function as agents in the world.

Acknowledgments

We thank Kajsa Igelstrom for comments. Research was supported by the Princeton Neuroscience Institute Innovation Fund.

Reprint requests should be sent to Michael S. A. Graziano, Department of Psychology, Peretsman-Scully Hall, Princeton University, Princeton, NJ 08544, or via e-mail: Graziano@princeton.edu.

REFERENCES

- Breitmeyer, B., & Ögmen, H. (2006). *Visual masking*. New York: Oxford University Press.
- Camacho, E. F., & Bordons Alba, C. (2004). *Model predictive control*. New York: Springer Publishing.
- Graziano, M. S. A. (2013). *Consciousness and the social brain*. New York: Oxford University Press.
- Graziano, M. S. A. (2014). Speculations on the evolution of awareness. *Journal of Cognitive Neuroscience*, 26, 1300–1304.
- Graziano, M. S. A., & Botvinick, M. M. (2002). How the brain represents the body: Insights from neurophysiology and psychology. In W. Prinz & B. Hommel (Eds.), *Common mechanisms in perception and action: Attention and performance XIX* (pp. 136–157). Oxford: Oxford University Press.
- Graziano, M. S. A., & Kastner, S. (2011). Human consciousness and its relationship to social neuroscience: A novel hypothesis. *Cognitive Neuroscience*, 2, 98–113.
- Graziano, M. S. A., & Webb, T. W. (2014). A mechanistic theory of consciousness. *International Journal of Machine Consciousness*, 6, 163–176.
- Hsieh, P., Colas, J. T., & Kanwisher, N. (2011). Unconscious pop-out: Attentional capture by unseen feature singletons only when top-down attention is available. *Psychological Science*, 22, 1220–1226.
- Ivanoff, J., & Klein, R. M. (2003). Orienting of attention without awareness is affected by measurement-induced attentional control settings. *Journal of Vision*, 3, 32–40.
- Jiang, Y., Costello, P., Fang, F., Huang, M., & He, S. (2006). A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of the National Academy of Sciences, U.S.A.*, 103, 17048–17052.
- Kelly, Y. T., Webb, T. W., Meier, J. D., Arcaro, M. J., & Graziano, M. S. A. (2014). Attributing awareness to oneself and to others. *Proceedings of the National Academy of Sciences, U.S.A.*, 111, 5012–5017.

- Kentridge, R. W., Nijboer, T. C., & Heywood, C. A. (2008). Attended but unseen: Visual attention is not sufficient for visual awareness. *Neuropsychologia*, *46*, 864–869.
- Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, *11*, 16–22.
- Lambert, A., Naikar, N., McLachlan, K., & Aitken, V. (1999). A new component of visual orienting: Implicit effects of peripheral information and subthreshold cues on covert attention. *Journal of Experimental Psychology—Human Perception and Performance*, *25*, 321–340.
- Lambert, A. J., Beard, C. T., & Thompson, R. J. (1988). Selective attention, visual laterality, awareness and perceiving the meaning of parafoveally presented words. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, *40A*, 615–652.
- Lamme, V. A. (2004). Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Networks*, *17*, 861–872.
- McCormick, P. A. (1997). Orienting attention without awareness. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 168–180.
- Norman, L. J., Heywood, C. A., & Kentridge, R. W. (2013). Object-based attention without awareness. *Psychological Science*, *24*, 836–843.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*, 3–25.
- Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. Bouwhuis (Eds.), *Attention and performance X* (pp. 531–556). London: Lawrence Erlbaum Assoc. Ltd.
- Posner, M. I., Rafal, R. D., Choate, L. S., & Vaughan, J. (1985). Inhibition of return: Neural basis and function. *Cognitive Neuropsychology*, *2*, 211–228.
- Scheidt, R. A., Conditt, M. A., Secco, E. L., & Mussa-Ivaldi, F. A. (2005). Interaction of visual and proprioceptive feedback during adaptation of human reaching movements. *Journal of Neurophysiology*, *93*, 3200–3213.
- Schurger, A., Pereira, F., Treisman, A., & Cohen, J. D. (2010). Reproducibility distinguishes conscious from nonconscious neural representations. *Science*, *327*, 97–99.
- Schurger, A., Sarigiannidis, I., Naccache, L., Sitt, J. D., & Dehaene, S. (2015). Cortical activity is more stable when sensory stimuli are consciously perceived. *Proceedings of the National Academy of Sciences, U.S.A.*, *112*, E2083–E2092.
- Tsushima, Y., Sasaki, Y., & Watanabe, T. (2006). Greater disruption due to failure of inhibitory control on an ambiguous distractor. *Science*, *314*, 1786–1788.
- Webb, T. W., & Graziano, M. S. A. (2015). The attention schema theory: A mechanistic account of subjective awareness. *Frontiers in Psychology*, *6*, 1–11.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, *269*, 1880–1882.