Attention, Consciousness, and Inattentional Blindness

By Alexander Wentzell[1]

NOTE This is a penultimate version of a forthcoming paper in Synthese. Please cite the published version.

Many philosophers propose an intimate link between attention and phenomenally conscious perception. This link comprises of two distinct theses. The first is that attention is necessary for phenomenally conscious perception. Accordingly, we consciously perceive only that to which we attend. Call this NEC. The second is that attention is sufficient for phenomenally conscious perception. Accordingly, we consciously perceive everything we attend. Call this SUFF. Call the family of views that endorse both NEC and SUFF *equivalency views*, or EVs, for short (Carruthers, 2015; De Brigard & Prinz, 2010; Posner, 1994; Prinz, 2012). If true, EVs would be a major boon to our theoretical understanding of phenomenal consciousness. In particular, it would greatly shape our understanding of how consciousness is neurologically implemented in the brain—something that remains largely mysterious. This is because EVs equate conscious states with attended states, and the neuroscience behind attended states is heavily studied.[2] Were EVs true, these would also be the mechanisms underlying our conscious experience of the world.

Contrary to EVs, several endorse what I'll call a dissociationist view of the relation between attention and phenomenally conscious perception. Dissociationist views may hold that there are interesting ways in which phenomenally conscious perception and attention interact. Nevertheless, they deny that attention is both necessary and sufficient for

[1] ACKNOWLEDGEMENTS: Thanks to Alexandra Cunningham and two anonymous referees for their helpful feedback on this manuscript.
[2] See, e.g., Lindsay (2020) and Kida and Okamoto (2023) for overviews.

phenomenally conscious perception. Thus, the dissociationist label captures a family of views; some deny that attention is necessary for phenomenally conscious perception (Smithies, 2011), some deny that it is sufficient (Baars, 1997; Cohen et al., 2012), and some deny both (Koch & Tsuchiya, 2007).

This paper argues for dissociationism. But rather than citing some yet unconsidered empirical phenomena to justify my claims, I think that the strongest argument in favor of dissociationism comes from an empirical phenomenon already widely discussed by proponents of EVs: inattentional blindness. Inattentional blindness (IB) occurs when subjects engaged in an attention-demanding task fail to notice an otherwise perfectly visible but task-irrelevant unexpected stimulus (Mack & Rock, 1998). IB is taken as a key piece of empirical evidence supporting EVs, specifically the necessity claim. After all, what inattentional blindness experiments purportedly demonstrate is just the necessity claim: when subjects fail to attend to the unexpected stimulus, they are not phenomenally conscious of it. Interpreted in this way, inattentional blindness experiments form a persuasive argument for attention being necessary for phenomenally conscious perception.

Here I argue that this thinking is backward: short of evincing the view, inattentional blindness experiments create a tricky dilemma for EVs. This is because, I argue, inattentional blindness experiments only provide support for NEC contingent on a certain interpretation of results, what I call the attentional patterns explanation (APE). According to APE, inattentional blindness occurs because of a difference in attentional patterns between those who see the unexpected stimulus and those who do not. Plainly, those who see the unexpected stimulus attend to it, and those who fail to see it do not. Considered in light of APE, inattentional blindness thus demonstrates NEC: without attention, subjects are not phenomenally conscious of the stimulus. In section 1, I explain this requirement in greater detail. In section 2, I argue that APE is contradicted by a wide range of empirical evidence. In section 3, I argue that the failure of APE leaves one of two interpretations of inattentional blindness experiments open. Depending on which of these two interpretations one favors, inattentional blindness experiments either 1) falsify NEC or 2) falsify SUFF and provide no empirical support for NEC. This creates a dilemma for proponents of EVs; either way they go, their theory is false. I close in section 4 by considering how the argument here also limits certain options in the dissociationist family of views.

Section 1: Inattentional Blindness and APE

"Consciousness" as a term admits many varieties. Phenomenally conscious mental states are typically understood as comprising of mental states for which there is "something it is like" for the subject to undergo (Nagel, 1974). Consider a subject experiencing a pain in their hand. In addition to the role this mental state plays in guiding the subject's behavior or factoring into their reasoning, there is also "something it is like" for the subject to experience this state—a first personal, qualitative feeling of pain, evanescent and difficult to describe but familiar to anyone who's undergone the state.

Phenomenally conscious mental states are typically contrasted with *access* conscious mental states. A mental state is access conscious if it is capable of (to some extent) factoring into a subject's reasoning, guidance of behavior, or verbal report (Block, 1995). Sometimes, a phenomenally conscious mental state will be access conscious, and vice-versa; consider our earlier example of a subject experiencing a pain in their hand. How these two types of consciousness are related beyond this, however, is a matter of ongoing debate. For some, all phenomenally conscious states are accessible, and vice-versa (Cohen & Dennett, 2011). For others, they may come apart: a mental state can be phenomenally conscious without being access conscious, and vice-versa (Block, 1995, 2007, 2011).

Attention, by contrast, is not a mental state but a mental process. We encounter, in most waking moments, a plurality of information. Attention refers to a subject's ability to prioritize and select some of this information over others in order to guide behavior, understood broadly to include both overt physical behaviors and mental actions like reasoning or imagining (James, 1890; Wu, 2024). Colloquial examples of attention are easy to generate. Imagine how you focus directly on the target when aiming a dart at a bullseye; other portions of your visual field recede into the background, and your increased focus on the bullseye allows you to better aim and throw the dart. As Wu (2024, p. 4-5) importantly notes, however, merely emphasizing the *selective* role of attention provides an inadequately robust definition. To borrow his example, cells in the retina engaged in sub-personal processing also select some information over others for further processing. What's important is *what type* of processing the selection feeds into. It's the hallmark of attention that this selection occurs for a specific type of processing, namely the guidance of subsequent *behavior* (Wu, 2024, p. 5-6).

The target reasoning of this paper takes attention to be necessary and sufficient for phenomenally conscious perceptual states. Framing the debate in terms of phenomenally conscious perception is a bit unconventional since most discussions frame the debate as being about whether attention is necessary and sufficient for phenomenal consciousness simpliciter. I have chosen to frame it this way for three reasons. First, visual perception is much more heavily studied and thus provides us with a richer set of reliable data in order to base our philosophical claims.[3] Second, the view that attention is necessary and sufficient for

[3] Thanks to an anonymous referee for suggesting this point.

phenomenal consciousness simpliciter entails the view that attention is necessary and sufficient for phenomenally conscious perception. Thus, if attention is not necessary and sufficient for phenomenally conscious perception, then it is not necessary and sufficient for phenomenal consciousness simpliciter. Third, this framing allows me to capture more nuanced views that maintain attention is necessary and/or sufficient for phenomenally conscious perception but deny that this is true of all forms of phenomenal consciousness. For example, Jennings (2020, p. 123) accepts NEC for phenomenally conscious perception. However, she rejects NEC for other forms of phenomenal consciousness (Jennings, 2015). With this clarified, I will now explain how inattentional blindness experiments purportedly provide empirical support for the target claim.

Inattentional blindness experiments are simple and elegant but produce quite surprising and unintuitive results. Experiments start with a primary task.[4] This varies depending on the experiment. In some, it involves multiple object tracking, in others visual search tasks, and in others length discrimination tasks. What they all have in common is that they require subjects to direct their attention toward a certain stimulus or set of stimuli to complete the task successfully.

Depending on the experiment, subjects may first engage in a few normal trials, or they may not. Regardless, on what is called the critical trial, the unexpected stimulus appears. The unexpected stimulus is absent from previous trials, distinct from other distractors in the display in some way, and perfectly visible to those not engaged in the primary task. After the critical trial, subjects are queried with increasingly leading questions to determine whether or not they saw the unexpected stimulus. Depending on their answer, subjects are subsequently

[4] For a detailed, book-length breakdown of this methodology, see Mack and Rock (1998).

coded as having either seen the unexpected stimulus or as having been inattentionally blind to it. After the critical trial, subjects undergo divided and full-attention trials, each of which is followed by an additional round of questions to determine whether they saw the stimulus. Divided-attention trials are just like the critical trial, with the only difference being that subjects now expect the unexpected stimulus to be present. In full-attention trials, subjects do not perform the task and merely watch the display. This is meant to control for the possibility that any failures to see the unexpected stimulus were due to the stimulus being somehow hidden in the display.

These experiments often find strikingly high rates of inattentional blindness during the critical trial. In one well-known experimental demonstration, 50% of observers tasked with counting passes made by a basketball team missed a person in a gorilla suit walking through the frame (Simons & Chabris, 1999). Likewise, 83% of radiologists examining CT scans of a lung failed to see a matchbook-sized picture of a gorilla embedded in the scan (Drew et al., 2013). And 65.5% of observers counting laps a swimmer took around a pool failed to notice an assault in the background (Cullen et al., 2022). Such errors are surprising but are also thought to have dramatic real-world consequences, like surgical errors (Hughes-Hallett et al., 2015) and vehicular accidents (Wang et al., 2022).

But what explains inattentional blindness? The intuitive answer is, as the name of the phenomenon suggests, *attention*. First, a bit of vocabulary: I will call how a subject distributes their attention over a range of stimuli over a certain time course their *attentional pattern*. I will reserve the locution "attends to X" to refer to a stimulus being selected by the particular attentional pattern realized by the subject. Subjects engaged in the primary task view the display with a certain attentional pattern to meet the demands of the task. To count the basketball passes, for example, they (primarily) attend to the basketball and not to other portions of the video. Due to this attentional pattern, subjects engaged in the task fail to attend to, and thus report seeing, the person in the gorilla suit. By contrast, those not engaged in the primary task view the display without such a focused attentional pattern. Without the need to meet task demands, their attention is more diffusely patterned around the display, and consequently, they attend to and report the person in the gorilla suit. This explanation is not off the cuff. It nicely explains why subjects *not* engaged in the primary task (or subjects in the full-attention trial) report seeing the stimulus. To summarize, then: what's the distinction between those who report the stimulus and those who don't? According to this line of interpretation, those who report the stimulus attend to it. Consequently, those who do not attend to the stimulus do not report it. Call this the "attentional patterns explanation," or APE for short.

APE makes no reference, however, to phenomenal consciousness. To provide support for NEC, it must be paired with a further premise: that subjects who do not report seeing the unexpected stimulus are not phenomenally conscious of it. This gives us the following argument for NEC. Call it "the APE Argument."

<u>The APE Argument:</u>

1) If subjects do not attend to X, then they do not report seeing X. (Premise/Attentional Patterns Explanation)
2) If subjects do not report seeing X, then subjects are not phenomenally conscious of X. (Further premise)
3) If subjects do not attend to X, then subjects aren't phenomenally conscious of X. (2,3, transitivity)

4) If subjects are phenomenally conscious of X, then subjects attend to X. In other words, attention is a necessary condition on phenomenal consciousness. (4, contraposition)

One argumentative strategy against the APE argument is to deny 2. In other words, one might deny that reports are accurate indicators of phenomenology. On such a view, subjects are still (at least to some extent) phenomenally conscious of the unexpected stimulus or its properties, but some just fail to report it (Kozuch, 2019). Inattentional blindness would then just demonstrate a failure of access consciousness (Block, 2011), and is perhaps better called “inattentional amnesia” (Moore, 2001; Wolfe, 1999). This would undermine NEC because subjects, according to this interpretation, are still phenomenally conscious of the unexpected stimulus. They just forget they saw it. Thus, even though they are not attending to the stimulus, they are still phenomenally conscious of it. And so, if this interpretation is correct, inattentional blindness experiments fail to demonstrate that attention is necessary for consciousness.

However, proponents of the APE argument have objected to this move (Suchy-Dicey, 2012). Furthermore, although it’s probably not possible to find definitive empirical evidence demonstrating that subjects are not phenomenally conscious of the unexpected stimulus when they don’t report it, there are some results that fit better with the interpretation implied by premise 2. One of these comes from an experiment by Becklen and Cervone (1983), which manipulated the gap between the disappearance of the unexpected stimulus on the critical trial and post trial questioning. For some subjects, they were queried immediately after the unexpected stimulus departed the display. For others, however, the task continued for an additional twenty-three seconds before questioning. If inattentional blindness were due to subjects forgetting they saw the stimulus, then one might expect a greater delay before post-trial questioning to result in lower reporting rates—since this greater delay would presumably allow the memory trace to get weaker. However, no difference was found.[5] Thus, for the sake of this essay, I will grant premise 2 of the argument.

Instead, I want to focus on premise 1. Premise 1, as stated, conflates two different groups of subjects in inattentional blindness experiments. In particular, it treats those participating in the primary task as a homogenous group, when this is not so. Crucially, *not all those engaged in the primary task are inattentionally blind to the unexpected stimulus*. What explains why these people, also engaged in the primary task, see the unexpected stimulus? To have a term, call only those subjects engaged in the task who fail to see the unexpected stimulus "non-seers," only those subjects engaged in the task who see the unexpected stimulus "task-seers," and subjects not engaged in the primary task "non-taskers." APE can be modified to incorporate this three-fold distinction by maintaining that the difference in attentional patterns it postulates between non-seers and non-taskers applies as well to non-seers and task-seers: Specifically, then, APE is committed to the conjunction of these two claims:

A. The attentional patterns of non-taskers and non-seers differ; Non-taskers attend to the unexpected stimulus, and thus report seeing it. Non-seers, due to having their attention taxed from the primary task, do not attend to and thus do not report seeing the unexpected stimulus.

[5] See also the discussion below about the lack of relation between inattentional blindness and working memory capacity.

B. The attentional patterns of task-seers and non-seers differ; although both groups have their attention taxed by the primary task, task-seers' attentional patterns are such that they attend to the unexpected stimulus, and thus report seeing it.

However, this modification poses a problem for APE. The reason why is, although the first conjunct of APE, (A), may very well be true, (B) is not. In the next section, I explain why.

Before getting there, however, it's important to note that, to support NEC, the proponent of the APE argument needs to be committed to both APE-(A) and APE-(B). To see why, imagine one adopted a modified form of APE that accepted APE-(A) but denied APE-(B). Then, APE will explain the different phenomenology between non-seers and non-taskers by appealing to attentional patterns but will deny that non-seers and task-seers attend differently. But, according to APE, non-seers do not attend to the unexpected stimulus. So, if the attentional patterns of non-seers and task-seers do not differ, then task-seers do not attend to the unexpected stimulus. But task-seers are phenomenally conscious of the stimulus, and we thus have a counterexample to NEC. So, to function as an argument for NEC, the proponent of APE is committed to the conjunction of (A) and (B).

Section 2: APE and Attentional Capture

In this section, I argue, on empirical grounds, that APE-(B) is false. This is because it leads to several empirical predictions which are unsupported given our current empirical evidence. APE explains IB in terms of a difference in attentional patterns leading to some subjects attending to the unexpected stimulus and some subjects not. In psychology, the unplanned orientation of attention towards a stimulus is referred to as *attentional capture* (Liesefeld et al., 2024; Theeuwes, 1992). Attentional capture is a familiar psychological phenomenon; consider the way a loud bang from across the room causes you to attend to its location. This is an instance of attentional capture: your attention is diverted toward a certain stimulus. From non-IB paradigms, we know quite a bit about attentional capture. We know which differences among individuals make one more susceptible to attentional capture. We know which stimuli make attentional capture more likely. And we know certain reliable behavioral indicators that attention has been captured. This yields three straightforward predictions that should all hold if APE-(B) is true:

1) Task-seers should demonstrate the behavioral indicators of attentional capture, and non-seers should not.
2) The individual differences that make one more prone to attentional capture should be found more in task-seers than in non-seers.
3) Those stimuli more likely to capture attention should be more likely to be attended to and thus seen at a higher rate.

All of these follow straightforwardly from the central role attentional capture plays in APE-(B). However, none of these three predictions are supported by our available empirical evidence. Let's consider each in turn.

2.1 Behavioural differences

Different attentional patterns result in certain noticeable behavioral changes. Thus, if task-seers are attending to the unexpected stimulus, we should expect them to show these noticeable behavioral changes. What changes to look for depends on the type of attention engaged, however. APE, as stated, is ambivalent, so we'll look at two possibilities. The first I'll survey is overt attention. Overt attention is the familiar form of visual attention that tracks one's gaze. Look one place, and look another, and you have just shifted your overt attention.

One obvious behavioral change that marks differences in overt attention is eye movements (Henderson, 2003; Theeuwes et al., 1998). After all, overt attention follows the gaze, and the gaze just is where the eyes are. One salient possibility, then, is that the difference in attentional patterns proposed by APE should result in a difference in gaze patterns. However, seers and non-seers are both equally likely to fixate on an unexpected stimulus and do so for similar durations (Beanland & Pammer, 2010; Memmert, 2006). So, the attentional difference proposed by APE-(B) cannot be overt.

The only remaining option is that APE-(B) postulate the difference in attention is *covert*. Covert attention refers to our ability to prioritize parts of the visual field away from fixation (Posner, 1980). It is both (perhaps) harder to grasp introspectively and harder to measure empirically. Could the proponent of APE-(B) maintain that the attentional difference between task-seers and non-seers is due to different *covert* attentional patterns?

First, it's worth noting that instances of overt attention—i.e., the movements of one's eyes—are not entirely disassociated from covert attentional patterns. Shifts in overt attention are preceded by shifts of covert attention, occurring approximately 100ms before eye movement (Peterson et al., 2004; Rizzolatti et al., 1987; Sheliga et al., 1994). This would indicate that non-seers *also* covertly attended to the unexpected stimulus.[6] Second, a difference in covert attentional patterns also produces certain behavioral differences, and these again aren't present between seers and non-seers. Particularly, allocations of covert attention are often measured using *task performance* (Folk et al., 1992; Posner, 1980). The reasoning proceeds as follows: if the addition of a task-irrelevant stimuli to the display (or: the

[6] Although note that the reverse need not be true—a shift in covert attention need not be followed by a shift in overt attention.

alteration of a feature of a distractor already present in the display) causes subjects' task performance to decrease, then the best explanation of this decrease in task performance is that the novel, task-irrelevant stimulus captured attention. Thus, if seers are covertly attending to the unexpected stimulus and non-seers are not, we should expect task-seers' performance to decrease.

Inattentional blindness studies often do not report accuracy rates of seers and non-seers, largely because, due to the fixed time course of inattentional blindness tasks, there is no worry of a potential speed/accuracy trade-off. However, when such numbers are reported, the overwhelming result is that no such difference is found (Beanland & Pammer, 2010; Bredemeier & Simons, 2012; Most et al., 2001; Redlich et al., 2019; Simons & Chabris, 1999). This again sits very uncomfortably with APE-(B). Were task-seers attending to the unexpecting stimulus and non-seers not, then the logic of attentional capture predicts that this should come at a performance cost.

A worry one might have is that a lack of difference in task performance is due to a ceiling effect.[7] If the task is easy, then a difference in attentional patterns might not result in a difference in task performance. This is because, due to the ease of the task, subjects might have leftover resources, allowing them to divert these surplus resources to the unexpected stimulus without impacting their ability to perform the task. However, a difference in task performance is not observed even when the primary task is attentionally demanding, precluding the possibility of a ceiling effect. To give one example, consider the task in experiment 1a of Beanland and Pammer (2010). This primary task involved subjects counting the number of bounces a group of target stimuli made of the edge of the display while ignoring

[7] Thanks to an anonymous reviewer for mentioning this possibility.

a distractor group. This task was difficult, however, with subjects underreporting the correct number of bounces by around 21%. However, the level of error was not significantly different between task-seers and non-seers (Beanland & Pammer, 2010, p. 980).[8]

2.2 Individual differences

Individual differences predict how prone one is to attentional capture (Hunt et al., 1989; Simons et al., 2024). APE explains inattentional blindness via attentional capture. In other words, those who see the unexpected stimulus have their attention captured by it. Thus, it's a natural thought that those subjects who possess individual differences making them more prone to capture should be more likely to see the unexpected stimulus than subjects who do not possess these individual differences. Several IB studies have been designed to this effect: first, subjects are tested for individual differences among a certain trait. Then, subjects perform an IB experiment. Experimenters then examine whether individual differences predict whether or not a subject sees the unexpected stimulus. If IB is due to attentional capture, as APE predicts, then those individual differences that make one more prone to attentional capture should also make one more likely to see the unexpected stimulus. There are several differences to look at. The first of these I will examine, in connection with IB, is *working memory capacity*.

[8] An anonymous reviewer makes two further points contesting this possibility: first, they point out that covert attention may still operate in a way that does not produce differences in task performance, if, for example, the primary task tapped into different resources. Second, they point out that different behavioural measures may indicate capture. Although I do not have space to fully adjudicate these matters, one interesting place we can look for further evidence is pupillometry. Pupil dilation has been shown to reliable track effort during attentionally demanding tasks like visual search (Porter et al., 2007). Consequently, one might expect those who display greater pupil dilation (and thus greater effort in the primary task) to be less likely to see the unexpected stimulation. However, pupil dilation, although it does track task accuracy, does not track seeing rates in inattentional blindness paradigms (Wright et al., 2013).

Working memory is a cognitive system that stores information in the service of ongoing tasks (Baddeley, 1992; Miller, 1960; Persuh et al., 2018). Working memory capacity varies amongst individuals and groups and can be measured using several different paradigms (Luck & Vogel, 2013; Wilhelm et al., 2013). WMC predicts performance on a large variety of cognitive tasks, but also certain attentional tasks like ignoring known distractor stimuli (Johnson et al., 2013; Unsworth & Engle, 2007). Indeed, WMC is also correlated with the ability to resist attentional capture (Lavie & De Fockert, 2005). Those with greater WMC are more resistant to distraction by task-irrelevant stimuli. We should thus expect those individuals with greater WMC to be *less* likely to have their attention captured by the unexpected stimulus and thus be less likely to report seeing it.

Original examinations of the relationship between WMC and IB found a positive correlation (Hannon & Richards, 2010; Richards et al., 2010). However, these findings were of borderline statistical significance (Beanland & Chan, 2016, p. 809), and have repeatedly failed to be replicated by subsequent studies (Beanland & Chan, 2016; Bredemeier & Simons, 2012; Kreitz, Furley, et al., 2015; Kreitz, Furley, Memmert, et al., 2016; Kreitz, Furley, Simons, et al., 2016). Working memory is related to one's susceptibility to attentional capture. But it doesn't seem to matter when it comes to inattentional blindness!

Another individual difference that has a predictable impact on attentional patterns is levels of emotional distress, particularly in terms of anxious arousal and anhedonic depression (Bredemeier et al., 2011; Levin et al., 2007). Interestingly, these pull in opposite ways, with anxious arousal being associated with increased vigilance and anhedonic depression being associated with lower susceptibility to attentional capture. Thus, individuals with greater anxiety should be less prone to inattentional blindness, and those with greater depression should be more prone. However, neither is related to individual blindness rates (Bredemeier et al., 2014).

Rather than looking for a variable that predicts attentional capture and then seeing if that variable predicts inattentional blindness rates, one can also directly compare an individual subject's performance in attentional capture tasks to their likelihood to see the unexpected stimulus. This was the strategy of Wright et al. (2018). First, they had subjects complete four traditional attentional capture paradigms. Subjects then completed a traditional inattentional blindness experiment. When comparing performance in the attentional capture paradigms to seeing rates in the IB task, Wright et al. found that performance in *none* of the four attentional capture paradigms was correlated with seeing rates in the inattentional blindness task.

As a recent meta-analysis by Simons et al. (2024) concludes, "[n]o individual difference measures reported in the published inattentional blindness literature strongly differentiated those who notice unexpected events from those who miss them." Given these individual differences *do* reliably predict attentional capture, this finding sits very uncomfortably with APE-(B).

2.3 Stimuli Type

Attentional capture can be divided into three distinct types. The first type of attentional capture is *bottom-up*. Bottom-up attentional capture is driven by features of the stimulus (Theeuwes, 2013). Accordingly, stimuli with certain properties, such as unique colors, abrupt onsets, or motion onsets generate a priority signal within the visual system and automatically capture the attention of the subject. Consider, for example, the way flashing emergency lights in your rearview mirror automatically capture your attention while driving.

The second type of attentional capture is *top-down.* Top-down attentional capture depends on the goals of the subject (Folk et al., 1992). Accordingly, the subject determines which features they wish to prioritize. These features form an “attentional set”, and stimuli with these features receive attentional prioritization. Consider, for example, looking for a book with a red spine among the stacks of a library. Those book spines matching this feature (e.g. those with red spines) cause you to automatically orient your attention to them, but this is contingent on you having set your goal to look for the red book. The last type of attentional capture is *value-based* (Anderson et al., 2011). Accordingly, stimuli that have previously been associated with reward generate a priority signal within the visual system and automatically capture the attention of the subject.

How these three types of attentional capture relate is an ongoing debate (Awh et al., 2012; Luck et al., 2021). One might view them all as distinct, or argue certain ones reduce to others. Regardless, each provides certain predictions about the type of stimuli that are more likely to capture attention, and thus the type of stimuli that should be more likely to be seen in an inattentional blindness paradigm.

To begin with bottom-up attentional capture: two particular types of stimuli are thought to capture attention in a bottom-up way. The first is color singletons, i.e. uniquely colored objects among a display (Belopolsky et al., 2010; Pashler, 1988). However, color singletons are not noticed more often in inattentional blindness tasks (Most et al., 2001). The second type of stimuli thought to reliably capture bottom-up attention is abrupt onsets, which do so more reliably than color singletons (Folk & Remington, 2015; Ruthruff et al., 2020; Zivony & Lamy, 2018). But again, in IB experiments, an unexpected stimulus with an abrupt onset is not seen more often than one that gradually enters the display (Most et al., 2005).

The types of stimuli that capture value-driven attention are more ad hoc. Subjects normally learn to associate a certain color dimension with a high reward value in an initial priming phase (Anderson et al., 2011). For example, subjects might receive a monetary reward every time they select red items from a display. On subsequent tasks, stimuli with the previously rewarded feature demonstrate greater attentional capture. However, previously rewarded stimuli are not noticed more often in inattentional blindness paradigms (Redlich et al., 2019). Another example of value-driven capture comes from food-related stimuli. Subjects who are hungry encounter greater attentional capture for food-related stimuli (Mogg et al., 1998; Piech et al., 2010). But again, despite being more likely to capture attention, food-related stimuli are not noticed more often by hungry subjects in inattentional blindness paradigms (Redlich et al., 2022).

Moving on to top-down attentional capture, this model predicts stimuli in line with the observer's current goals (what's called their *attentional set*) will capture attention. An endlessly replicated result amongst inattentional blindness studies is that seeing rates increase the more similar the unexpected stimulus is to the expected, task-relevant stimuli. This is true both in terms of properties like shape and color (Most et al., 2005) but also in terms of more abstract similarities like semantic category (Koivisto & Revonsuo, 2007; Most, 2013). *Prima facie*, this may seem like very good evidence that IB is related to top-down attention. But closer examination reveals that the way categorical similarity modulates seeing in inattentional blindness does differ from how it modulates attentional patterns via top-down mechanisms.

When subjects form a top-down attention set for a particular semantic category, their attention is consequently direct toward not just objects of that category, but objects that are highly associated with it. For example, when given a visual search task where they have to see if a monkey is present among a display, subjects will direct their attention toward items that are associated with monkeys, like bananas, at a rate higher than non-associated distractors (Moores et al., 2003). Thus, were inattentional blindness due to subjects selectively deploying their attention toward a certain attentional-set, an unexpected stimulus that is semantically associated with the task-relevant stimuli should also capture attention and thus be more likely to be seen. In order to test this connection, Clement et al. (2019) performed the following experiments. Subjects viewed a display containing two groups of four images. One group was pictures of monkeys, and the other rabbits. Subjects were tasked with either tracking the monkeys or the rabbits and counting the number of bounces pictures of this type made off the edge of the display. They ran two different experiments that differed in what they used as an unexpected stimulus. In the first, the unexpected stimulus was either a novel monkey or rabbit. Predictably, when the unexpected stimulus was a monkey, subjects tracking the monkeys saw the unexpected stimulus far more often than those tracking the rabbit, and when the unexpected stimulus was a rabbit, subjects tracking the rabbits saw the unexpected stimulus far more often than those tracking the monkeys. In the second experiment, however, the unexpected stimulus was semantically associated with the task relevant stimuli. On some trials, it was a banana, which is associated with monkeys, and on others a carrot, which is associated with rabbits. Unlike when the unexpected stimulus shared a semantic category with the task-relevant stimuli, subjects were *not* more likely to notice an object when it was associated with the task relevant stimuli. But this is in tension with what we know about top-down attention. If inattentional blindness were due to subjects selectively attending to stimuli in line with their attentional set, then they should be more likely to attend to and thus report seeing semantically associated stimuli. But no such increase in seeing rates were found.

One might be tempted to turn this reasoning around: if semantic associations do not modulate awareness in inattentional blindness paradigms, then maybe we should take this as evidence that top-down attentional sets do not also bias attention toward semantically related stimuli. My response is that, when considered in conjunction with the other distinctions between capture and blindness, as a package, this indicates that we should be skeptical that it is an attentional effect. What we've observed is not just one distinction between how attentional capture functions and inattentional blindness, but several.

2.4: Where we stand

APE yields several empirical predictions that I've surveyed here, none of which are supported by our available empirical evidence. It's worth cautioning at this point that this conclusion is based on null results—or studies *failing* to find an effect rather than establishing that an effect is present. Null results are often considered less informative, because they can happen for many reasons, such as unideal experimental design or poor statistical power. The possibility cannot be ruled out that with better methods or more highly powered studies, the effects entailed by APE may turn up. But while this is an empirical possibility, what we have here is not just a single null result, but a robust and established pattern of null results across converging predictions and methods. Were just one of these results true, then I think drawing any sort of conclusions about APE would be premature. But I hope that I've presented several converging lines of evidence which indicate the following: current empirical evidence does not support the assertion that there is *any* difference in attentional patterns between task-seers and non-seers. This is bad news for APE. In the next sections, I outline the consequences this has for EVs.

Section 3: APE and EVs, Redux

This ends my discussion of APE. One might think that the falsity of APE just denies EVs (or rather, one of its conjuncts) one form of empirical support. But here I want to argue that the falsity of APE-(B) leads to a dilemma that, depending on how it is solved, provides either a counterexample to NEC or SUFF, thus falsifying EVs.

The dilemma is generated as follows: empirical evidence indicates that the attentional patterns of task-seers and non-seers do not differ. In other words, there are no relevant differences in how they attend to the stimuli in the display. This claim is consistent with two distinct possibilities: either (D1) their attentional patterns are such that they attend to the unexpected stimulus, or (D2) their attentional patterns are such that they do not attend to the unexpected stimulus.

Recall too another commitment that proponents of EVs made to get the argument from APE off the ground: that the lack of report from non-seers indicates a lack of phenomenology. This commitment, combined with the fact that either (D1) or (D2) above must be the case, means that either NEC or SUFF is false. Here's how:

Say that (D1) is true. Then, by stipulation, non-seers attend to the unexpected stimulus. But, given the argument from APE's commitment to reports being accurate indicators of phenomenology, non-seers are not phenomenally conscious of the unexpected stimulus. Thus, attention to a stimulus is not sufficient for phenomenally conscious perception. SUFF is false.

Say that (D2) is true. Then, by stipulation, non-seers *do not* attend to the unexpected stimulus. The empirical evidence surveyed in the previous section indicates that there is no relevant way in which the attentional patterns of non-seers and task-seers differ. Thus, it follows that task-seers *also do not* attend to the unexpected stimulus. But, given they report it, task-seers *are* phenomenally conscious of the unexpected stimulus. Thus, attention to a stimulus is not necessary for phenomenally conscious perception. NEC is false.

One horn of the dilemma has to be true. EVs thus ought to be rejected. Dissociationism follows.

Section 4: IB and Dissociationism

The preceding argument supports a dissociationist view about the relation between attention and phenomenally conscious perception. But it also has implications for *which* dissociationist view is right. Recall that dissociationism, as I've defined it, refers to a family of views encompassing both views that deny either NEC or SUFF, and views that deny both. We might call views that deny either NEC or SUFF but not both "single dissociation views" and views that deny both NEC and SUFF "double dissociation views."

My sympathies lie with the double dissociationist view. However, arguing this fully is outside the scope of this paper. The argument presented here does not quite get us to settling the question of whether a double or single dissociation view is correct. But it does, however, have some bearing on *which* single dissociation view is defendable. This is because, depending on which horn of the disjunct is true, either NEC is false or SUFF is false. And, if NEC is false, then the corresponding single dissociation theory that denies only SUFF is false, and vice-versa.

Figuring out which horn of the disjunct—either (D1) or (D2)—is true is a complicated task. How do we tell if non-seers and task-seers attend to the stimulus? In traditional attention paradigms, attention to a stimulus is measured using task performance, but this requires a control—a group of subjects completing the task for whom we can confidently say attention is *not* directed towards the unexpected stimulus. There does not appear to be any such control available to us. Other strategies are needed.

One option is to cash out "attending to a stimulus" in terms of something *other* than task performance. An example of this comes from Wu (2014)'s "empirically sufficient condition." according to which selection of a stimulus for a task is sufficient for attention to that stimulus. Task-seers select the unexpected stimulus for a task, namely that of reporting its presence. Thus, according to Wu's empirically sufficient condition, task-seers attend to the unexpected stimulus. And, because, as I have argued, the attentional patterns of task-seers and non-seers do not differ, it thus follows that non-seers attend to the stimulus. So (D1) is correct, and the sufficiency version of the single endorsement view is false.

But I do not think this reasoning is ideal. It rules out the possibility of non-seers attending to the unexpected stimulus by fiat. By EVs's lights, the lack of phenomenal consciousness is what explains the lack of report in non-seers. If we identify a lack of attention with a lack of report, then it's conceptually impossible for non-seers to attend to the unexpected stimulus. But whether or not they are attending to the unexpected stimulus seems like an empirical matter, not something that's conceptually impossible. An empirically sufficient condition that disentangles attending from the performance or non-performance of a certain task would be better.

One alternative approach is to treat subjects' performance on pre-critical trials as a control and look for any difference in performance between these trials and critical trials. After all, subjects can't attend to something that isn't there! This data, however, is often unreported in IB experiments. In one study I'm aware of that does explicitly mention this comparison (Most et al., 2005), they found a decrease in counting accuracy *both* among non-seers and task-seers on the critical trial compared to previous, non-critical trials. This would indicate attentional capture by the unexpected stimulus for both non-seers and task-seers. Thus, this line of reasoning again suggests (D1) is true, and SUFF is false. However, this strategy for assessing whether the unexpected stimulus was attended to is again imperfect. By design, the critical trial contains an additional stimulus, not present in the previous trials. This leaves open the possibility that any decrease in performance is caused simply by the increased processing demands of viewing a more complex scene, rather than attentional capture.

Regardless of their imperfections, here are two strands of evidence that suggest resolving the disjunct such that SUFF is false. On the face of it, this may seem like a happy result for the proponents of NEC. However, this is not so. This is because inattentional blindness is one of the most ubiquitously cited behavioral demonstrations of NEC. But, by diffusing the dilemma such that SUFF is false, proponents of NEC must interpret inattentional blindness experiments in a way such that they no longer provide *any* empirical support for the necessity claim. Here is how:

Recall *why* inattentional blindness was taken to demonstrate the necessity claim: When experimental conditions are such that attention is *not* directed toward the unexpected stimulus, non-seers fail to be phenomenally conscious of it. As a phenomenon, it demonstrates that, in the absence of attention, consciousness is lost. Thus, the thought goes, attention is necessary for phenomenal consciousness.

But resolving the dilemma in a way that doesn't falsify NEC requires abandoning this explanation. It requires accepting the claim that *both non-seers and task-seers attend to the unexpected stimulus*—i.e., accepting (D1). Thus, it is *not* true that inattentional blindness experiments provide any support for the necessity claim because inattentional blindness experiments *do not* manipulate attention away from the unexpected stimulus. Therefore, although resolving the dilemma in this way does not falsify NEC, it does deny proponents of NEC one of the most celebrated forms of empirical support their view has.

CONCLUSION: So What's the mechanism?

The discussion thus far has left one question unaddressed: if APE does not explain inattentional blindness, then what does? What's the mechanism that explains why some people see the target stimuli, and others do not? As an anonymous reviewer points out, this question being left open is not just an intriguing consequence of my argument here but instead may jeopardize the whole analysis. After all, APE is a widely accepted, intuitive explanation of why inattentional blindness happens. As such, one might be tempted to conclude from the empirical results not that APE is false, but that the methods we use to behaviorally measure attentional patterns are instead flawed or insufficient. This is especially so if there exists no plausible alternative mechanism to attention that can explain these results.[9]

The first thing to say in response to this worry, is that one might still be hesitant, even without a plausible alternative mechanism, to draw this conclusion. Attention is one of the most widely studied phenomena in cognitive psychology. And, in other tasks where attention is plausibly invoked, such as visual search, the methods relied on here (stimuli changes, individual differences, task performance) have been reliably used to track attention for several decades (Wolfe, 2018). And so, we should be hesitant to throw the baby out with the bathwater.

[9] Thanks to an anonymous reviewer for raising this point.

Nevertheless, it is certainly true that my argument here is constrained by our current methods of studying and tracking how attention is distributed across a scene. While I maintain that the results surveyed do rely on our best methods of behaviorally tracking attention, it would be naïve to deny that someday more sensitive measures of detecting attentional differences may be developed, and that these may vindicate APE. As such, the conclusion of this paper is perhaps best understood as a disjunction—either there is no difference between the attentional patterns of seers and non-seers, as current evidence suggests, or our current means of measuring attention are insufficiently granular to pick up on such a difference.[10]

However, in absence of novel ways of measuring attention, I do think there are two other explanations for what's going on in IB that do not centralize the role of attention as APE does. The first of these is mentioned earlier: that IB is better explained as a failure of memory. According to this view, non-seers are perceptually similar to task-seers; both form a perceptual representation of the unexpected stimulus. But they differ in whether or not this representation gets encoded into memory. Task-seers successfully encode the representation in memory, and non-seers don't. This explains why they differ in their reporting of the unexpected stimulus, without appealing to a difference in attention.

In light of the discussion here, this alternative proposal deserves consideration, but as mentioned previously it's not without faults. I, too, do not favor this explanation. This stems from another empirical finding mentioned earlier: null results regarding working

[10] An anonymous referee raises the possibility of gist perception as an alternative mechanism to explain APE. Gist perception refers to subject's ability to rapidly extract statistical information and semantic categories from scenes (Potter et al., 2014) Two worries one might have of this alternative are as follows. First, in some prominent models of attention, recovery of gist occurs pre-attentively (Wolfe, 2021). Second, the sorts of properties that are recovered in gist perception concern groups of objects, not individuals. However, this possibility merits further consideration.

memory capacity and inattentional blindness. If IB occurred due to a failure to encode representations of the unexpected stimulus in memory, then one would expect subjects with greater working memory capacity to be more likely to retain, and thus report seeing the unexpected stimulus. But this result has not been vindicated by empirical research. An anonymous reviewer suggests another hurdle for this view to overcome: it needs to explain *why* non-seers fail to encode the unexpected stimulus in working memory. This seems particularly pressing for this interpretation, given the close relation between attention and working memory encoding (Oberauer, 2019). In short, attending to a stimulus often results in better encoding. And so, a natural explanation of *why* one group forgets the unexpected stimulus is because the unexpected stimulus was attended to by seers/not attended to by non-seers.[11] But this is what APE claims, and this view is meant to be an alternative to APE. Nevertheless, given the failure of APE, this alternative proposal should be considered a viable option, and more work needs to be done to examine whether it can overcome the flaws identified.

The second alternative proposal I'll mention stems from a confound in the design of inattentional blindness experiments. In order to manipulate what subjects attend to, experimenters provide them with task instructions, under the innocuous assumption that

[11] One possible route a proponent of this view may take is to appeal to discussions about depth of processing (Craik & Tulving, 1975). This literature identifies several features of a task that result in better encoding and subsequent retrieval of stimuli at test. To give one example, asking subjects if a word fits into a sentence results in better encoding for a later test than asking subjects if a word is upper case. Plausibly, these effects operate independently of attention. This is because more complex and (presumably) attention demanding tasks do not result in better encoding if they only probe "shallow" features of the stimulus (like its distribution of vowels/consonants, see Craik & Tulving (1975), experiment 5). One such encoding effect is the congruity effect (Schulman, 1974). Accordingly, stimuli judged congruent with a cue are better remembered than those that are incongruent (e.g., if the cue is "is this word uppercase?" words that *are* presented in uppercase will be better recalled at test). Potentially, a bridge could be made here to inattentional blindness experiments, where the main finding of note is that recall is better when stimuli are congruent with the target category. But more work needs to be done by proponents of this explanation to flesh out this potentiality, especially since most of the encoding work cited concerns words.

subjects will translate these instructions into a particular attentional pattern. However, I think it's plausible that task instructions, in addition to causing the subject to approach the task with a certain attentional pattern, also cause them to form certain expectations regarding what will be present in the display. If, for example, you tell subjects that their goal is to count the passes made by the white shirt team, then they will form the expectation that the display will contain a team of basketball players with white shirts. An unstated assumption in the inattentional blindness literature is that it's the first effect of this manipulation, the alteration of attentional patterns, that's causally responsible for the results. To the extent that expectations play any causal role, it's only mediately through their impact of how the subject distributes their attention. But, given the evidence surveyed here, perhaps it's the second manipulation that's the culprit: in short, inattentional blindness occurs not due to a subject's attentional patterns impacting what they see, but due to subjects' expectations regarding what will be present in the display directly impacting what they see. Indeed, the confound of manipulating attention and manipulating expectations was realized by early researchers on IB. In their seminal work on the topic, Mack and Rock (1998, p. 204) themselves recognized that "failures to perceive could be due to either the irrelevant, attention-demanding task, to the absence of expectation or intention to the critical stimulus, or both." Similarly, a response to the book by Jochen Braun (2001) develops this suggestion further. I'll call this proposal *the expectation interpretation.*

Again, the proponent of APE need not deny that inattentional blindness experiments manipulate subject's expectations in the way identified. It is fully consistent with APE to acquiesce to this point. They can maintain that the causal role of expectations is *mediate*—it's only because expectations prime or facilitate attentional processes that they have any role in the phenomenon. But what the proponent of APE must deny—and what the expectation interpretation maintains—is that there is a direct causal impact of subject's expectations upon their perceptual processes, operating independently of attention.[12] This view owes us three things. First, it must conceptually distinguish between attentional states and expectational states. Second, it must motivate the idea that expectational states can, independently of attention, impact perceptual processing. And third, it must plausibly demonstrate that something like this is going on in inattentional blindness experiments.

I maintain that the many difficulties faced by APE lend some prima facie plausibility to the expectations interpretation. Although I leave a more nuanced discussion of its merits to future work, I do want to close by highlighting how the view can make progress in meeting these three desiderata. The best way to empirically ground this view comes from work in cognitive neuroscience that distinguishes between attention and expectations in perception (Aitken et al., 2020; Summerfield & Egner, 2016; Summerfield & Egner, 2009). This line of research distinguishes attention and expectation as follows. Attending to a stimulus X involves, as we've seen, prioritizing the processing of X in service of task demands. Expecting a stimulus X, by contrast, involves forming a prior that the occurrence or presence of X is likely in the sensory environment. Conceptually, these types of states are doubly-dissociable. First, you can expect to see X without attending to X. As you search a bookshelf for your red spined book, you expect to see several books, but you only attend to a subset of them in the course of your search. Second, you can attend to X without expecting to see X. This seems to be what happens in most cases of "bottom-up" capture, where a stimulus captures attention due to, for example, an abrupt onset into the scene.

[12] Accepting this view need not entail accepting that there's no causal interface between expectations and attention.

One way this distinction is introduced into experimental settings is by presenting subjects with a predictive cue to a feature of an oncoming stimulus, and then varying whether that cued-feature is task-relevant or irrelevant. For example, in Kok et al. (2012), subjects were presented, sequentially, two gratings that differed in orientation, spatial frequency, and contrast. Each grating was viewed for 550ms, with a 100ms space in between. Prior to the stimulus onset, an auditory cue played indicating (with 75% accuracy) whether the grating was left or right oriented. In the first block of trials, subjects' task involved determining whether the grating of the second stimulus was clockwise or counter-clockwise to the first. Thus, in this block, the predictive cue provides task relevant information that subjects can incorporate into their attentional set. In the second block, subjects' task is to determine whether the second grating is of higher contrast than the first. Thus, given it is task irrelevant, and given the short presentation time of the stimulus, the cue does not provide task-relevant/attentional information to the subjects—although it does provide information regarding what they should expect to see in the display.

Interestingly, the two different types of information seem to impact neural processing in different ways. Attention operates, neurally, by boosting the response of neurons that process the particular attended dimension, relative to when these dimensions are unattended. As Maunsell (2015) writes,

> It has long been known that shifts of attention toward or away from a visual stimulus are associated with changes in the way that neurons respond to that stimulus. These include modulations of the latency, magnitude and variability of neuronal responses. Such changes in neuronal responses are most easily described when a single stimulus is presented in a neuron's receptive field and attention is directed toward or away

> from that stimulus. In that case, attention to a receptive field stimulus is usually associated with responses that are faster, stronger, and less variable compared to responses when attention is directed elsewhere.

Expectations, by contrast, have a dampening impact on neural activity (Alink et al., 2010; Kok, Jehee, et al., 2012). When a stimulus is expected, it produces reduced neural response to when it's unexpected, demonstrating the opposite neural signal as attention (Kok, Rahnev, et al., 2012). Perhaps the most natural way to interpret this finding is that expected stimuli are less represented by neural activity. This has been dupped the "expectation suppression effect" in the literature (Feuerriegel et al., 2021).[13] But another possibility is that this decrease in neural response is due to what's called a sharpening effect (Kok, Jehee, et al., 2012). Because the stimuli type is expected, subject's perceptual systems are able to represent it more efficiently. Interestingly, the study by Kok et a. (2012) discussed earlier found evidence for this interpretation. A classifier that looked at neural responses and tried to predict the orientation of the viewed gratings was much more accurate when examining neural responses to expected rather than unexpected stimuli.

The neuroscience underlying this distinction is largely still nascent, but I hope to have illustrated that this alternative enjoys some empirical plausibility from non-IB literature. Now, let's discuss the third desiderata: can we plausibly make the case that something like this is going on in inattentional blindness experiments? Well, the expectation interpretation fits comfortably with by far the most widely replicated result in the IB literature: that subjects are more likely to see an unexpected stimulus more similar to the expected

[13] A worry one might have is that this effect is better explained as the unexpected stimulus capturing attention, rather than a decrease in processing to the expected stimulus (Rungratsameetaweemana & Serences, 2019). However, it's worth nothing that there's no reason to think the unexpected stimulus would capture attention, given it's equally salient to the task relevant information.

stimulus present in the display. Consider, again, Simmons and Charbris (1999) gorilla experiment. Here, subjects attend to one of the two basketballs being passed by either the black shirts or the white shirts. Given they are tracking a ball being passed by members of one team, they form the expectation that they'll see, over the course of the experiment, that color. And, it's because of this antecedent expectation that subject's tracking the black shirts report seeing the gorilla more often—because they expect to see black stimuli in the display.

Assuming the underlying distinction between expectations and attention is valid, then, it is challenging to see how one can cleanly disentangle the two in an inattentional blindness experiment. One attempt pursued by Mack and Rock (1998, p. 209) involved having a primary task that is very attentionally undemanding. In their experiment, they had subjects merely view a display for 5 seconds under the pretense of exploring afterimages—thus, subjects only real task involved attending to a fixation cross. Regardless of the lack of primary task demand, they still found an inattentional blindness rate of 44%, which they took as "fairly solid evidence that the lack of expectation alone, without the additional factor of attention to another task, can cause IB and poor shape perception" (Mack & Rock, 1998, p. 209).[14]

Although I think this view is certainly worth pursuing, it's not without its faults: in particular, subjects within conditions (e.g., subjects who are tracking the black shirt players) presumably approach the display with the same set of expectations; however, they nevertheless differ in whether or not they see the unexpected stimulus. For the proponent of the expectations interpretation, there are a couple of responses available. They may

[14] Of course, this doesn't completely eliminate attention and is far from a perfect demonstration.

maintain that the influence of expectations on perception is stochastic or probabilistic rather than determinate. Or, they may maintain that whether expectations impact perception as stated is determined by further factors of the cognitive agent that moderate, but do not mediate, the effect. Interestingly, this second line of response enjoys some level of empirical support in studies that examine susceptibility of inattentional blindness and personality factors. A large study in this area found susceptibility to inattentional blindness highly correlated with *openness*, a personality characteristic that captures one's receptiveness to new experiences (Kreitz, Schnuerch, et al., 2015). As the authors of that study hypothesize, "it is conceivable that individuals who are open to new experiences might also be more receptive (in a perceptual sense) to unexpected objects or events" (Kreitz, Schnuerch, et al., 2015, p. 3).[15] However, further experimental validation is still warranted. Although my sketch of it here is largely gestural, I hope it provides yet another serious contender when it comes to what the mechanism behind inattentional blindness, if not attention, might be.

[15] This problem actually equally faces APE as well: APE maintains that different attentional patterns are responsible for the difference in seeing, but presumably those subjects within the same condition approach the display with similar attentional patterns. As an anonymous reviewer points out, this consideration could equally be used by proponents of APE in order to answer this concern, since it's plausible something like openness would also impact attention. But this runs into the same problem as earlier: then why is there no difference in traditional attentional measures between seers and non-seers?

References

Aitken, F., Menelaou, G., Warrington, O., Koolschijn, R. S., Corbin, N., Callaghan, M. F., & Kok, P. (2020). Prior expectations evoke stimulus-specific activity in the deep layers of the primary visual cortex. *PLOS Biology*, *18*(12), e3001023. https://doi.org/10.1371/journal.pbio.3001023

Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W., & Muckli, L. (2010). Stimulus Predictability Reduces Responses in Primary Visual Cortex. *The Journal of Neuroscience*, *30*(8), 2960–2966. https://doi.org/10.1523/JNEUROSCI.3730-10.2010

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences*, *108*(25), 10367–10371. https://doi.org/10.1073/pnas.1104047108

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443. https://doi.org/10.1016/j.tics.2012.06.010

Baars, B. J. (1997). Some Essential Differences between Consciousness and Attention, Perception, and Working Memory. *Consciousness and Cognition*, *6*(2–3), 363–371. https://doi.org/10.1006/ccog.1997.0307

Baddeley, A. (1992). Working Memory. *Science*, *255*(5044), 556–559. https://doi.org/10.1126/science.1736359

Beanland, V., & Chan, E. H. C. (2016). The relationship between sustained inattentional blindness and working memory capacity. *Attention, Perception & Psychophysics*, *78*(3), 808–817. https://doi.org/10.3758/s13414-015-1027-x

Beanland, V., & Pammer, K. (2010). Looking without seeing or seeing without looking? Eye movements in sustained inattentional blindness. *Vision Research*, *50*(10), 977–988. https://doi.org/10.1016/j.visres.2010.02.024

Becklen, R., & Cervone, D. (1983). Selective looking and the noticing of unexpected events. *Memory & Cognition*, *11*(6), 601–608. https://doi.org/10.3758/BF03198284

Belopolsky, A. V., Schreij, D., & Theeuwes, J. (2010). What is top-down about contingent capture? *Attention, Perception & Psychophysics*, *72*(2), 326–341. https://doi.org/10.3758/APP.72.2.326

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, *18*, 227–287.

Block, N. (2007). Consciousness, Accessibility, and the Mesh Between Psychology and Neuroscience. *Behavioral and Brain Sciences*, *30*(5), 481–548. https://doi.org/10.1017/s0140525x07002786

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, *15*(12), 567–575. https://doi.org/10.1016/j.tics.2011.11.001

Braun, J. (2001). It's Great But Not Necessarily About Attention. *Psyche*, *7*(6).

Bredemeier, K., Berenbaum, H., Most, S. B., & Simons, D. J. (2011). Links between neuroticism, emotional distress, and disengaging attention: Evidence from a single-target RSVP task. *Cognition & Emotion*, *25*(8), 1510–1519. https://doi.org/10.1080/02699931.2010.549460

Bredemeier, K., Hur, J., Berenbaum, H., Heller, W., & Simons, D. J. (2014). Individual differences in emotional distress and susceptibility to inattentional blindness. *Psychology of Consciousness: Theory, Research, and Practice*, *1*(4), 370–386. https://doi.org/10.1037/cns0000032

Bredemeier, K., & Simons, D. J. (2012). Working memory and inattentional blindness. *Psychonomic Bulletin & Review*, *19*(2), 239–244. https://doi.org/10.3758/s13423-011-0204-8

Carruthers, P. (2015). *The Centered Mind: What the Science of Working Memory Shows Us About the Nature of Human Thought*. Oxford University Press.

Clement, A., Stothart, C., Drew, T., & Brockmole, J. R. (2019). Semantic associations do not modulate the visual awareness of objects. *Quarterly Journal of Experimental Psychology*, *72*(5), 1224–1232. https://doi.org/10.1177/1747021818811411

Cohen, M. A., Cavanagh, P., Chun, M. M., & Nakayama, K. (2012). The attentional requirements of consciousness. *Trends in Cognitive Sciences*, *16*(8), 411–417. https://doi.org/10.1016/j.tics.2012.06.013

Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, *15*(8), 358–364. https://doi.org/10.1016/j.tics.2011.06.008

Craik, F. I., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology. General*, *104*(3), 268–294. https://doi.org/10.1037/0096-3445.104.3.268

Cullen, H. J., Paterson, H. M., & van Golde, C. (2022). Does experiencing inattentional blindness for crime influence eyewitness recall? *Memory*, *30*(2), 206–216. https://doi.org/10.1080/09658211.2021.2002906

De Brigard, F., & Prinz, J. (2010). Attention and consciousness. *WIREs Cognitive Science*, *1*(1), 51–59. https://doi.org/10.1002/wcs.27

Drew, T., Võ, M. L.-H., & Wolfe, J. M. (2013). The Invisible Gorilla Strikes Again: Sustained Inattentional Blindness in Expert Observers. *Psychological Science*, *24*(9), 1848–1853.

Feuerriegel, D., Vogels, R., & Kovács, G. (2021). Evaluating the evidence for expectation suppression in the visual system. *Neuroscience & Biobehavioral Reviews*, *126*, 368–381. https://doi.org/10.1016/j.neubiorev.2021.04.002

Folk, C. L., & Remington, R. W. (2015). Unexpected Abrupt Onsets Can Override a Top-Down Set for Color. *Journal of Experimental Psychology. Human Perception and Performance*, *41*(4), 1153–1165. https://doi.org/10.1037/xhp0000084

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(4), 1030–1044. https://doi.org/10.1037/0096-1523.18.4.1030

Hannon, E. M., & Richards, A. (2010). Is Inattentional Blindness Related to Individual Differences in Visual Working Memory Capacity or Executive Control Functioning? *Perception*, *39*(3), 309–319. https://doi.org/10.1068/p6379

Henderson, J. M. (2003). Human gaze control during real-world scene perception. *Trends in Cognitive Sciences*, *7*(11), 498–504. https://doi.org/10.1016/j.tics.2003.09.006

Hughes-Hallett, A., Mayer, E. K., Marcus, H. J., Pratt, P., Mason, S., Darzi, A. W., & Vale, J. A. (2015). Inattention blindness in surgery. *Surgical Endoscopy*, *29*(11), 3184–3189. https://doi.org/10.1007/s00464-014-4051-3

Hunt, E., Pellegrino, J. W., & Yee, P. L. (1989). Individual Differences in Attention. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 24, pp. 285–310). Academic Press. https://doi.org/10.1016/S0079-7421(08)60540-X

James, W. (1890). *The Principles of Psychology* (Vol. 1). H. Holt and Company.

Jennings, C. D. (2015). Consciousness Without Attention. *Journal of the American Philosophical Association*, *1*(2), 276–295. https://doi.org/10.1017/apa.2014.14

Jennings, C. D. (2020). *The Attending Mind* (New edition). Cambridge University Press.

Johnson, M. K., McMahon, R. P., Robinson, B. M., Harvey, A. N., Hahn, B., Leonard, C. J., Luck, S. J., & Gold, J. M. (2013). The relationship between working memory capacity and broad measures of cognitive ability in healthy adults and people with schizophrenia. *Neuropsychology*, *27*(2), 220–229. https://doi.org/10.1037/a0032060

Kida, T., & Okamoto, H. (2023). Editorial: New insights into the cognitive neuroscience of attention. *Frontiers in Human Neuroscience*, *17*. https://www.frontiersin.org/articles/10.3389/fnhum.2023.1338738

Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, *11*(1), 16–22. https://doi.org/10.1016/j.tics.2006.10.012

Koivisto, M., & Revonsuo, A. (2007). How Meaning Shapes Seeing. *Psychological Science*, *18*(10), 845–849. https://doi.org/10.1111/j.1467-9280.2007.01989.x

Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, *75*(2), 265–270. https://doi.org/10.1016/j.neuron.2012.04.034

Kok, P., Rahnev, D., Jehee, J. F. M., Lau, H. C., & De Lange, F. P. (2012). Attention Reverses the Effect of Prediction in Silencing Sensory Signals. *Cerebral Cortex*, *22*(9), 2197–2206. https://doi.org/10.1093/cercor/bhr310

Kozuch, B. (2019). Gorillas in the Missed (but Not the Unseen): Reevaluating the Evidence for Attention Being Necessary for Consciousness. *Mind and Language*, *34*(3), 299–316. https://doi.org/10.1111/mila.12216

Kreitz, C., Furley, P., Memmert, D., & Simons, D. J. (2015). Inattentional Blindness and Individual Differences in Cognitive Abilities. *PLOS ONE*, *10*(8), e0134675. https://doi.org/10.1371/journal.pone.0134675

Kreitz, C., Furley, P., Memmert, D., & Simons, D. J. (2016). The Influence of Attention Set, Working Memory Capacity, and Expectations on Inattentional Blindness. *Perception*, *45*(4), 386–399. https://doi.org/10.1177/0301006615614465

Kreitz, C., Furley, P., Simons, D. J., & Memmert, D. (2016). Does working memory capacity predict cross-modally induced failures of awareness? *Consciousness and Cognition*, *39*, 18–27. https://doi.org/10.1016/j.concog.2015.11.010

Kreitz, C., Schnuerch, R., Gibbons, H., & Memmert, D. (2015). Some See It, Some Don’t: Exploring the Relation between Inattentional Blindness and Personality Factors. *PLOS ONE*, *10*(5), e0128158. https://doi.org/10.1371/journal.pone.0128158

Lavie, N., & De Fockert, J. (2005). The role of working memory in attentional capture. *Psychonomic Bulletin & Review*, *12*(4), 669–674. https://doi.org/10.3758/BF03196756

Levin, R. L., Heller, W., Mohanty, A., Herrington, J. D., & Miller, G. A. (2007). Cognitive deficits in depression and functional specificity of regional brain activity. *Cognitive Therapy and Research*, *31*(2), 211–233. https://doi.org/10.1007/s10608-007-9128-z

Liesefeld, H. R., Lamy, D., Gaspelin, N., Geng, J. J., Kerzel, D., Schall, J. D., Allen, H. A., Anderson, B. A., Boettcher, S., Busch, N. A., Carlisle, N. B., Colonius, H., Draschkow, D., Egeth, H., Leber, A. B., Müller, H. J., Röer, J. P., Schubö, A., Slagter, H. A., … Wolfe, J. (2024). Terms of debate: Consensus definitions to guide the scientific discourse on visual distraction. *Attention, Perception, & Psychophysics*, *86*(5), 1445–1472. https://doi.org/10.3758/s13414-023-02820-3

Lindsay, G. W. (2020). Attention in Psychology, Neuroscience, and Machine Learning. *Frontiers in Computational Neuroscience*, *14*. https://www.frontiersin.org/articles/10.3389/fncom.2020.00029

Luck, S. J., Gaspelin, N., Folk, C. L., Remington, R. W., & Theeuwes, J. (2021). Progress Toward Resolving the Attentional Capture Debate. *Visual Cognition*, *29*(1), 1–21. https://doi.org/10.1080/13506285.2020.1848949

Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: From psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, *17*(8), 391–400. https://doi.org/10.1016/j.tics.2013.06.006

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

Maunsell, J. H. R. (2015). Neuronal Mechanisms of Visual Attention. *Annual Review of Vision Science*, *1*(1), 373–391. https://doi.org/10.1146/annurev-vision-082114-035431

Memmert, D. (2006). The effects of eye movements, age, and expertise on inattentional blindness. *Consciousness and Cognition*, *15*(3), 620–627. https://doi.org/10.1016/j.concog.2006.01.001

Miller, G. A. (1960). *Plans and the structure of behavior*. Holt.

Mogg, K., Bradley, B. P., Hyare, H., & Lee, S. (1998). Selective attention to food-related stimuli in hunger: Are attentional biases specific to emotional and psychopathological states, or are they also found in normal drive states? *Behaviour Research and Therapy*, *36*(2), 227–237. https://doi.org/10.1016/S0005-7967(97)00062-4

Moore, C. M. (2001). Inattentional blindness: Perception or memory and what does it matter? *Psyche: An Interdisciplinary Journal of Research on Consciousness*, *7*(2), No Pagination Specified-No Pagination Specified.

Moores, E., Laiti, L., & Chelazzi, L. (2003). Associative knowledge controls deployment of visual selective attention. *Nature Neuroscience*, *6*(2), 182–189. https://doi.org/10.1038/nn996

Most, S. B. (2013). Setting sights higher: Category-level attentional set modulates sustained inattentional blindness. *Psychological Research*, *77*(2), 139–146. https://doi.org/10.1007/s00426-011-0379-7

Most, S. B., Scholl, B. J., Clifford, E. R., & Simons, D. J. (2005). What You See Is What You Set: Sustained Inattentional Blindness and the Capture of Awareness. *Psychological Review*, *112*(1), 217–242. https://doi.org/10.1037/0033-295X.112.1.217

Most, S. B., Simons, D. J., Scholl, B. J., Jimenez, R., Clifford, E., & Chabris, C. F. (2001). How Not to Be Seen: The Contribution of Similarity and Selective Ignoring to Sustained Inattentional Blindness. *Psychological Science*, *12*(1), 9–17. https://doi.org/10.1111/1467-9280.00303

Nagel, T. (1974). What Is It Like to Be a Bat? *The Philosophical Review*, *83*(4), 435. https://doi.org/10.2307/2183914

Oberauer, K. (2019). Working Memory and Attention – A Conceptual Analysis and Review. *Journal of Cognition*, *2*(1), 36. https://doi.org/10.5334/joc.58

Pashler, H. (1988). Cross-dimensional interaction and texture segregation. *Perception & Psychophysics*, *43*(4), 307–318. https://doi.org/10.3758/BF03208800

Persuh, M., LaRock, E., & Berger, J. (2018). Working Memory and Consciousness: The Current State of Play. *Frontiers in Human Neuroscience*, *12*. https://www.frontiersin.org/articles/10.3389/fnhum.2018.00078

Peterson, M. S., Kramer, A. F., & Irwin, D. E. (2004). Covert shifts of attention precede involuntary eye movements. *Perception & Psychophysics*, *66*(3), 398–405. https://doi.org/10.3758/bf03194888

Piech, R. M., Pastorino, M. T., & Zald, D. H. (2010). All I saw was the cake. Hunger effects on attentional capture by visual food cues. *Appetite*, *54*(3), 579–582. https://doi.org/10.1016/j.appet.2009.11.003

Porter, G., Troscianko, T., & Gilchrist, I. D. (2007). Effort during visual search and counting: Insights from pupillometry. *Quarterly Journal of Experimental Psychology*, *60*(2), 211–229. https://doi.org/10.1080/17470210600673818

Posner, M. I. (1980). Orienting of Attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25. https://doi.org/10.1080/00335558008248231

Posner, M. I. (1994). Attention: The mechanisms of consciousness. *Proceedings of the National Academy of Sciences*, *91*(16), 7398–7403. https://doi.org/10.1073/pnas.91.16.7398

Potter, M. C., Wyble, B., Hagmann, C. E., & McCourt, E. S. (2014). Detecting meaning in RSVP at 13 ms per picture. *Attention, Perception, & Psychophysics*, *76*(2), 270–279. https://doi.org/10.3758/s13414-013-0605-z

Prinz, J. (2012). *The Conscious Brain*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195314595.001.0001

Redlich, D., Memmert, D., & Kreitz, C. (2022). Does hunger promote the detection of foods? The effect of value on inattentional blindness. *Psychological Research*, *86*(1), 98–109. https://doi.org/10.1007/s00426-021-01480-y

Redlich, D., Schnuerch, R., Memmert, D., & Kreitz, C. (2019). Dollars do not determine detection: Monetary value associated with unexpected objects does not affect the likelihood of inattentional blindness. *Quarterly Journal of Experimental Psychology*, *72*(9), 2141–2154. https://doi.org/10.1177/1747021819835148

Richards, A., Hannon, E. M., & Derakshan, N. (2010). Predicting and manipulating the incidence of inattentional blindness. *Psychological Research*, *74*(6), 513–523. https://doi.org/10.1007/s00426-009-0273-8

Rizzolatti, G., Riggio, L., Dascola, I., & Umiltá, C. (1987). Reorienting attention across the horizontal and vertical meridians: Evidence in favor of a premotor theory of attention. *Neuropsychologia*, *25*(1A), 31–40. https://doi.org/10.1016/0028-3932(87)90041-8

Rungratsameetaweemana, N., & Serences, J. T. (2019). Dissociating the impact of attention and expectation on early sensory processing. *Current Opinion in Psychology*, *29*, 181–186. https://doi.org/10.1016/j.copsyc.2019.03.014

Ruthruff, E., Faulks, M., Maxwell, J. W., & Gaspelin, N. (2020). Attentional dwelling and capture by color singletons. *Attention, Perception & Psychophysics*, *82*(6), 3048–3064. https://doi.org/10.3758/s13414-020-02054-7

Schulman, A. I. (1974). Memory for words recently classified. *Memory & Cognition*, *2*(1), 47–52. https://doi.org/10.3758/BF03197491

Sheliga, B. M., Riggio, L., & Rizzolatti, G. (1994). Orienting of attention and eye movements. *Experimental Brain Research*, *98*(3), 507–522. https://doi.org/10.1007/BF00233988

Simons, D. J., & Chabris, C. F. (1999). Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events. *Perception*, *28*(9), 1059–1074. https://doi.org/10.1068/p281059

Simons, D. J., Hults, C. M., & Ding, Y. (2024). Individual differences in inattentional blindness. *Psychonomic Bulletin & Review*. https://doi.org/10.3758/s13423-023-02431-x

Smithies, D. (2011). Attention Is Rational-Access Consciousness. In C. Mole, W. Wu, & D. Smithies (Eds.), *Attention: Philosophical and psychological essays* (pp. 247–276). Oxford University Press.

Suchy-Dicey, C. M. (2012). Inductive parsimony and the Methodological Argument. *Consciousness and Cognition*, *21*(2), 605–609. https://doi.org/10.1016/j.concog.2010.11.011

Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, *13*(9), 403–409.

Summerfield, C., & Egner, T. (2016). Feature-based attention and feature-based expectation. *Trends in Cognitive Sciences*, *20*(6), 401–404. https://doi.org/10.1016/j.tics.2016.03.008

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, *51*(6), 599–606. https://doi.org/10.3758/BF03211656

Theeuwes, J. (2013). Feature-based attention: It is all bottom-up priming. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *368*(1628), 20130055. https://doi.org/10.1098/rstb.2013.0055

Theeuwes, J., Kramer, A. F., Hahn, S., & Irwin, D. E. (1998). Our Eyes Do Not Always Go Where We Want Them to Go: Capture of the Eyes by New Objects. *Psychological Science*, *9*(5), 379–385. https://doi.org/10.1111/1467-9280.00071

Unsworth, N., & Engle, R. W. (2007). The Nature of Individual Differences in Working Memory Capacity: Active Maintenance in Primary Memory and Controlled Search From Secondary Memory. *Psychological Review*, *114*(1), 104–132. https://doi.org/10.1037/0033-295X.114.1.104

Wang, Y., Wu, Y., Chen, C., Wu, B., Ma, S., Wang, D., Li, H., & Yang, Z. (2022). Inattentional Blindness in Augmented Reality Head-Up Display-Assisted Driving. *International Journal of Human–Computer Interaction*, *38*(9), 837–850. https://doi.org/10.1080/10447318.2021.1970434

Wilhelm, O., Hildebrandt, A., & Oberauer, K. (2013). What is working memory capacity, and how can we measure it? *Frontiers in Psychology*, *4*. https://www.frontiersin.org/articles/10.3389/fpsyg.2013.00433

Wolfe, J. M. (1999). Inattentional Amnesia. In V. Coltheart (Ed.), *Fleeting memories cognition of brief visual stimuli*. MIT Press.

http://JE5QH2YG7P.search.serialssolutions.com/?V=1.0&L=JE5QH2YG7P&S=JCs&C=TC0000155954&T=marc

Wolfe, J. M. (2018). Visual Search. In J. T. Wixted (Ed.), *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience* (1st ed., pp. 569–623). Wiley. https://doi.org/10.1002/9781119170174

Wolfe, J. M. (2021). Guided Search 6.0: An updated model of visual search. *Psychonomic Bulletin & Review*, *28*(4), 1060–1092. https://doi.org/10.3758/s13423-020-01859-9

Wright, T. J., Boot, W. R., & Morgan, C. S. (2013). Pupillary response predicts multiple object tracking load, error rate, and conscientiousness, but not inattentional blindness. *Acta Psychologica*, *144*(1), 6–11. https://doi.org/10.1016/j.actpsy.2013.04.018

Wright, T. J., Roque, N. A., Boot, W. R., & Stothart, C. (2018). Attention Capture, Processing Speed, and Inattentional Blindness. *Acta Psychologica*, *190*, 72–77. https://doi.org/10.1016/j.actpsy.2018.07.005

Wu, W. (2014). *Attention* (1st edition). Routledge.

Wu, W. (2024). We know what attention is! *Trends in Cognitive Sciences*, *28*(4), 304–318. https://doi.org/10.1016/j.tics.2023.11.007

Zivony, A., & Lamy, D. (2018). Contingent Attentional Engagement: Stimulus- and Goal-Driven Capture Have Qualitatively Different Consequences. *Psychological Science*, *29*(12), 1930–1941. https://doi.org/10.1177/0956797618799302