

Distracted Drivers and Unattended Experience

Wayne Wright

wrightwt@mindstuff.net

Consider the much-discussed case of the distracted driver, who is alleged to successfully navigate his car for miles despite being completely oblivious to his visual states. Perhaps he is deeply engrossed in the music playing over the radio or in philosophical reflection, and as a result he goes about unaware of the scene unfolding before him on the road. That the distracted driver has visual experiences of which he is not aware is a possibility that first-order representationalists (FOR) happily accept, but higher-order representationalists (HOR) steadfastly deny. HOR claims that perceptual states become conscious only as the object of higher-order states; perceptual states are not intrinsically conscious. According to HOR, since the driver is supposed to be completely distracted by other cognitive tasks, he cannot form higher-order representations of his visual states, with the result that those states are disqualified as experiences.¹ HOR theories have come in two flavors,¹ those that claim that the relevant higher-order representations are thought-like (HOT) and those that rely on an inner perception-like mechanism that is directed toward one's mental states (HOP).² Central to FOR is the idea that if a perceptual state has the appropriate pedigree - typically involving its relations to other first-order cognitive states - it is conscious, independent of whether the perceiver forms a higher-order representation of it. In view of the markedly different conclusions HOR and FOR reach about it, the distracted driver scenario presents a good example of the sort of issues these positions are in dispute over.

In this note, I use data from empirical research on attention – particularly work on inattention blindness and change blindness, but also studies of the cognitive demands of driving - to argue that we should reject both the standard FOR interpretation of the distracted driver case and the possibility of unattended experience altogether. I aim to do this in a way that results in the endorsement of a form of higher-order theory that is largely consistent with the core claims of a specific version of FOR (Michael Tye’s PANIC theory), but which adds the requirement that conscious perceptual states have attentional resources devoted to them. Contributing a further category to this debate, I will refer to my view as a form of higher-order attention theory (HOA). The ultimate goal of this reinterpretation of the distracted driver case and focus on empirical research on attention is to move representationalism (and theories of the mind in general) away from a mistaken “inner snapshot” view of experience and the empirically dubious accounts of experience it gives rise to.

1. Details of the distracted driver case

In its simplest detail, the distracted driver case consists of a subject S having driven his car for some time and suddenly coming to the realization that he is not at all aware of what he has been doing behind the wheel up to that moment. S has no recollection of changing gears, stopping at red lights, making steering adjustments to avoid potholes, and so forth. S also realizes that he has been deeply engrossed in some activity other than driving. Up to this point, I see nothing to object to, as the greater attention S has been paying to, say, the concert broadcast on his radio would likely swamp memory of the details of any visual episodes he might have had during his

drive. I suppose that most readers have found themselves in a similar situation and, much to my dread, I have occasionally had the same shocking realization myself.

The next step is far more controversial. Both FOR and HOR theorists agree that, while distracted, S is lacking higher-order consciousness of his visual states; S does not have thoughts about or inner perceptual episodes directed toward his visual states. According to HOR, this is sufficient for determining that the driver had no consciousness at all of his visual states while he was distracted; one becomes conscious of mental states only by forming higher-order representations of them. The FOR theorist, however, marks a distinction between higher-order consciousness and phenomenal consciousness. This enables the FOR theorist to grant that S lacks higher-order consciousness of his visual states while maintaining that there is a very real and important sense in which S is conscious of his visual states. The driver is unaware of how the road before him looks, but it still looks some way to him; he is still having visual experiences of the road.

There is a deeper issue about the distracted driver case that has thus far been overlooked, one that, once considered, should prove to be of great benefit in settling the dispute between HOR and FOR. If the distracted driver is truly “oblivious to [his] experiences” and is not paying any attention to them (Tye 1995, 115), the case seems implausible on empirical grounds. The attentional requirements of driving are such that, once considered, they force a revision of the distracted driver case into something far weaker than the claim that the driver is subject to visual states that he is unaware of the entire time he is driving. The significance of this is that the same reasons that lead us to doubt that the driver could successfully navigate his car in the complete absence of attention should also prompt us to question the empirical possibility of experience without attention. Carrying further the worries raised by studies of the cognitive demands of

driving, other research on the role of attention in our visual experience indicates that there are good empirical motivations for doubting the FOR theorist's contention that one can experience in the absence of attention.

The considerable difficulties FOR faces in meeting its empirical commitments should lead us to endorse some form of HOR. I will explain my resulting HOR view in greater detail as the paper progresses, but for now it will suffice to say that I do not take the higher-order states necessary for consciousness to be akin to an internal form of perception, to be thought-like, or to themselves be conscious. The position I am ultimately led to is a form of HOR that takes the relevant higher-order states to be attentional. Since the idea that inner sense or inner perception is what makes first-order mental states conscious invites the suspicion that a Cartesian materialist view of experience is on offer and the attentional mechanisms my account draws on do not have the same vulnerability, I take my HOA position to be distinct from HOP.³

In what follows, I will present three reasons for questioning the standard FOR position on unconscious experience. The first two are presented in the next section and are targeted at the distracted driver case itself; one more of a methodological concern about how the case is presented and interpreted and the other an empirical point based on studies of the cognitive demands of driving. In some sense, both points also apply to typical HOR theories, since it seems as though most HOR theorists have been willing to grant the FOR theorist's description of the scenario while disagreeing with the FOR theorist's conclusions about whether the driver was subject to visual experiences. The third point, which has broader consequences for the prospects of FOR, is presented in the paper's third section and is based on research on inattentional blindness and change blindness, and the consequences of that research for the idea that there is unattended experience. Before concluding with brief comments that describe what looks to be

the best way to go about developing the HOA theory I want to advance, I address some possible objections to my view.

2. Concerns about oblivious driving

A worry about the distracted driver case that immediately comes to mind is that it seems as though the case is vastly underdescribed to support the strong conclusion that S has been unaware all along of his visual states. Something seems wrong with the inference from the fact that S suddenly realizes he is unaware of what has been going on up to a certain point to the claim that S was completely unaware of his visual states while they took place. A standard way that the case has appeared in the philosophical literature moves quickly from the subject's astonishment that he has no idea what he has been doing and his realization that his mind has been preoccupied with some other cognitive task besides driving, to the conclusion that the subject was unaware of his visual states (and, presumably, his driving behavior) the entire time.

Consider the following from David Armstrong:

After driving for long periods of time, particularly at night, it is possible to “come to” and realize that for some time past one has been driving without being aware of what one has been doing. The coming-to is an alarming experience. It is natural to describe what went on before one came to by saying that during that time, one lacked consciousness.⁴

(Armstrong 1980, 59)

Another telling of the distracted driver story, this one from Michael Tye, takes as a given that the driver was totally unaware of his visual states and driving behavior the entire time he was piloting his vehicle:

[the distracted driver] is thinking hard about philosophy, say, as he drives along the highway. The driver is unaware of how the road ahead looks to him, of the visual experiences he is undergoing. But the experiences are still there. He still sees the road ahead. (Tye 2000, 36. His emphasis)

It should be evident that, if we stick to its basic outline, there is an at least equally plausible rival explanation of the distracted driver case that is left unexplored in these presentations of the scenario. That explanation is that S was aware of the visual scene before him for at least part of the time and he used the visual features he attended to in guiding his driving behavior, but the far greater share of cognitive resources devoted to some other task prevented S from storing any detail about his experience. Since S is dedicating the bulk of his cognitive resources to another task, it should come as no surprise to find that information about the visual scene is quickly bumped from working memory and is never encoded in long-term memory; the appropriate encoding conditions fail to obtain due to lack of available cognitive resources. Thus, at times, S could have been attending to features of the scene before his eyes, engaged in appropriate driving behavior on the basis of whatever he was (minimally) aware of, and have no memory at all of having done so. This would satisfactorily explain both S's success at keeping his car on the road and his shock when he comes to.

The second problem is that the distracted driver case, as it is typically elaborated, is not even empirically plausible, simply given what we know about the cognitive demands of driving. A totally distracted driver, one who is oblivious to his visual sensations of the road before him, would not be able to successfully navigate the road for very long at all, much less for the extended periods sometimes claimed in presentations of the distracted driver case. Consider the

following descriptions of the cognitive demands of driving, the first two from US studies and the third from a Canadian study:

The task of driving involves a complicated interaction of psychological, physical, cognitive, psychomotor, and sensory skills, placing high attentional demands on drivers. (Anderson, et al 2002, 6)

As the driving environment becomes more complex, more attentional resources are needed to maintain performance.... Importantly, attentional demand, as measured by mean reaction time, was strongly related to the density of accidents along the routes [of the study]. (Lee, et al 2000, 3)

Driving is a complex behavior that requires the extraction and integration of information from multiple sources in an effort to produce safe and efficient vehicle control.... Much of the information relevant to driving is taken in visually; consequently any change in drivers' visual behavior could be significant for driving safety. (Harbluk et al 2002, 3).

Not surprisingly, all three studies note that as drivers increasingly devote their attention to other tasks (e.g., cell phone conversations, visually scanning the vehicle's gauges, reflection on mathematical puzzles), they became increasingly distracted from the task of driving. Even non-visual driving-oriented tasks can interfere with the cognitive processing required to support safe driving; e.g., speech-interaction with onboard systems. There is only so much attention to go around. Although simple driving environments such as highways (usually the cited environment in presentations of the distracted driver case) are known to have a lower attentional demand for safe driving than complex environments (e.g., the cramped, windy roads of a Swiss village),

there seems to be no indication that safe driving can take place in any environment with an extended, complete absence of awareness of the road and one's driving behavior. The conclusion to draw from the empirical evidence is that distracted drivers, if they are successful in navigating the road for any significant length of time, have at least some awareness, although perhaps quite insubstantial and fleeting, of the visual scene. Without sufficient attention being paid to one's visual experience and driving behavior, one will quickly find one's car quite mangled.

The bit of reasoning just employed runs squarely opposite to that of the FOR theorist.

Consider the following passage:

... the fact that I am sometimes distracted by philosophical thoughts when I drive [does not entail] that I no longer see the road and the cars ahead. I am able at such times to attend to my visual sensations even though I do not do so. But the visual sensations are there all right. How else do I keep the car on the road? (Tye 2000, 14)

Tye takes it as a given that the distracted driver is not attending to his experiences and pairs that assumption with the driver's success in navigating the road to come to the conclusion that the driver must nonetheless be subject to visual experiences of the road. I urge an entirely different way of approaching things. The empirical evidence indicates that if one is successfully navigating the road, one is paying sufficient attention to one's visual states. Since the distracted driver manages to keep his car on the road for a considerable time, what we should instead conclude about the distracted driver is that he is the subject of visual states of which he has at least some minimal awareness.

One further aspect of Tye's position to consider is his claim that he is "able at such times to attend to [his] visual sensations even though [he does] not do so." Depending on how the

ability is to be understood, this is not so obviously true. In a trivial sense, of course the distracted driver is able to attend to certain features of his visual states that he in fact is not paying attention to. All he would have to do is re-direct his attention to the currently unattended features of the visual scene and he would be attending to them. That is, perceiver P is not attending to feature F of visual scene V at time₀, but P is attending to F of V at time₁. There is also a counterfactual understanding of such an ability that is of no assistance to the FOR theorist; viz., P did not attend to F of V at time₀, but in different circumstances, P would have attended to F of V at time₀. However, it seems as though Tye is claiming – and the FOR theorist requires – something stronger. What is at stake is whether the visual information not attended to is cognitively available to P while it is not attended to. That is, whether the information about F of V acquired through an unattended visual state at time₀ is potentially available for P to use in thought and the direction of intentional action, even though P does not actually do so. If it turns out that that information is not cognitively available in the way just described, then there is no reason to take P's visual state at time₀ as a visual experience. This is so by Tye's own lights, as it would violate the very plausible condition of poise that he contends is essential to phenomenally conscious states.⁵ As I will discuss in the next section, work on change blindness and inattention blindness reveals that a lack of attention does more than simply block belief formation on the basis of one's visual states; i.e., the information is present, but use of it is interfered with because one's attention is elsewhere. The unattended information is not even potentially available for one to use in thought and direction of intentional action. This has damaging consequences for the FOR theorist's position.

Returning to the studies on driving, their results show that not only are attentional and processing resources compromised by multi-tasking while driving, but that there is an overall

reduction in perceptual activity. Saccades (rapid, targeted eye movements used to explore the visual environment) significantly decrease and “visual tunneling” (i.e., focusing on the central area of the visual field to the exclusion of the areas to the right and to the left) dramatically increases in response to increased cognitive demand for other tasks (Harbluk et al 2002, 9-12). This reduction in perceptual activity can plausibly be interpreted in a way that further explains why one might have the sense that one has been driving around completely unaware of the visual scene, despite that not actually being the case. Not only are attentional and memory resources compromised by the other cognitive tasks the distracted driver is engaged in, but there is a vast reduction in information acquisition about the visual scene (when compared to attentive driving). The highly distracted driver is having what we might regard as very dim experiences that receive greatly reduced attention. Thus he is picking up very little information about the scene before his eyes and that information acquisition is taking place in cognitive conditions that are far less than optimal for recording in memory. It is no wonder that, when he emerges from his mental fog, it seems to the distracted driver as though he had been driving in an unconscious state the entire time; there is no information about his recent driving episode to reflect on to convince him otherwise.

My claim is that the distracted driver is subject to visual states that are accompanied by enough awareness, a sufficient amount of attentional resources, to enable him to keep the car on the road, shift gears when required, stop at red lights and so forth. Of course, the driver’s ability to carry out such tasks and to respond to emergency situations is greatly compromised by the lack of attention, but if there were an utter lack of attention to the visual scene, if the driver were completely oblivious of his visual states in the way the FOR theorist has typically contended, the consequences would be quite grave for the distracted driver; he would soon run his car into some

sort of hazard! What is important to note is that this treatment of the distracted driver case does not demand that the distracted driver have higher-order thoughts about his visual states in order to be conscious of them. Nor should the relevant resources be thought of as a form of inner perception of one's experiences. However, once the need for attention is brought into the picture, the view becomes a HOR theory, since "attention is inherently intentional ... [it] must be directed to some thing" (Mack and Rock 1998a, 4. Their emphasis). These higher-order attentional states are not conscious (unless they, too, become the object of a higher-order state), but it is the attachment of attentional resources to sensory outputs (of the right sort) that makes a state conscious. In the case of a visual experience, it is the world-directed visual state itself that is conscious, becoming so by having attentional resources devoted to it.

I should be careful here to clarify the kind of attention I am drawing on in my treatment of the distracted driver case, since 'attention' turns out to refer to a collection of heterogeneous cognitive processes. For example, the kind of attentional mechanisms implicated in the account of consciousness I am advocating are not akin to the subconscious mechanisms described by Anne Treisman in accounting for the binding of primitive representations of features into object representations.⁶ The attentional resources I have in mind – and which seem to be at work in the earlier cited studies of the cognitive demands of driving – operate at a conscious level and have to do with access to and control over information by the perceiving subject.⁷ The relevant attentional mechanisms mark certain features of the visual scene as salient, indexing them for further searches and processing. Attention comes in degrees, and the less attentive one is to one's perceptual activity, the less control over and access to perceptual information one has. What I have on offer is a form of HOR that does not place such demanding cognitive requirements on the perceiver that it eliminates as potential subjects of experiences creatures

(such as human infants and dogs) that it is prima facie compelling to think have actually experiences. In this regard HOA likely has more in common with FOR than it does with HOT.⁸

HOA does not deny the possibility of information pickup through visual stimulation that we are not aware of, but that plays a role in guiding our behavior. In fact, research has shown that unconsciously acquired information can be used to guide behavior even when it conflicts with the information currently in attention.⁹ I am instead arguing two other points. The first is that it is a mistake to suppose that a creature can engage in as complex a behavior as driving for an extended period of time without attending to its relevant perceptual states. Empirical data support this claim and also provide the resources needed to account for the compelling, albeit erroneous, impression that the driver is completely oblivious to his visual states during his entire driving episode.

The second point, which is strongly gestured at by the new look at the distracted driver case and which will be directly addressed in the next section, is that the kind of perceptual sensitivity that is to count as experience is of a sort that is accompanied by attention. Not just any form of information acquisition through the eyes should qualify as visual experience, as a state accompanied by a distinctive phenomenal character and available to make an impact on the belief/desire system. While experience is obviously useful for directing simple tasks, it is compelling to think that some such tasks could be executed on the basis of unconsciously acquired visual information. Engagement in complex tasks, the sort of tasks that require integration of information from multiple sources and sophisticated behavioral outputs, demands a certain level of attentional resources for successful completion. More complex tasks also look to be ideal candidates for the kind of behavior that requires visual experience – not just visual stimulation - as an input. Certainly, it is natural to conceive of experience as a state capable of

making a noticeable difference in our deliberations and guidance of behavior, whereas “mere visual stimulation” is ill-suited to play such a robust role in our cognitive economies. This intersection of the attentional demands of complex behaviors and the role of experience as an input to the reasoning processes that result in complex behaviors should lead us to take seriously the idea that attention and experience are connected in a deep and important way.

We can see how this distinction between simple and complex tasks, and the attentional requirements of each, fits the case of the distracted driver. If the driver is truly oblivious, he likely could pilot the car a short time without incident, in the right settings, solely on the basis of information acquired through unattended visual states. This is especially so if the driver is driving in a familiar or simple environment. I will concede that there probably are stretches during an extended period of distracted driving in which the driver is truly oblivious. However, the distracted driver’s continued success at keeping the car on the road for an extended time, accelerating and decelerating, shifting gears when appropriate, and so forth looks to be explainable only by assuming that he was paying some degree of attention to his visual states during at least some points during his drive; that he had control over and access to information about the visual scene and his driving activity. I suspect that whatever surface plausibility is attached to the FOR theorist’s contention that the driver is completely unaware of his visual experiences the entire drive derives from focus on the likelihood of the driver having moments of actual unconsciousness regarding his visual states to the exclusion of the requirement that those oblivious moments be interspersed with spells of (at least) minimal attention paid to his informationally impoverished visual states. In the next section, I will present further empirical reasons for thinking that we do not have cognitive access to unattended features of the visual scene and thus that we experience only that to which we attend.

3. Attention and visual states

Results from studies of inattention blindness (IB) and change blindness (CB) indicate that the claim that subjects can have visual experiences that they are not aware of is wrong.¹⁰ CB is the failure to see unattended changes in a visual scene. Due to a distracting disturbance, a perceiver fails to notice large-scale changes in the scene directly before his eyes, changes that are in full view and which would be easily detected in circumstances in which the perceiver's attention was not drawn elsewhere. IB is a stronger sort of failure to perceive what is directly before one's eyes; research on IB reveals that perceivers do not consciously perceive features of the visual scene that they do not attend to. I say "stronger" because CB is concerned with transitions of features across perceptual episodes, whereas the target phenomena of IB are stimuli actually present in an isolated occurrent perceptual episode. From the point of view of our current explanatory interests, IB makes a more imposing case against the possibility of unconscious experience than CB does.

The upshot of this empirical research is that the idea that we experience a far greater wealth of detail than we actually attend to, the impression that "all of the information in the visual environment is potentially available for attentive processing,"¹¹ is mistaken; we might say that visual experience turns out to be a good deal less than what meets the eye. I term this faulty view of experience the 'snapshot' or 'inner picture' view, since it treats visual experiences as internal picture-like representations that can be scanned at will by the perceiver.¹² Contrary to the snapshot view, those features of our occurrent visual episodes that we do not attend to are not conscious in any genuine way. They are "lost" once not attended to and thus are unable to make

any impact on the belief/desire system, also cutting off their possible links with intentional control of action. However, the dubious snapshot view of experience is clearly what the FOR theorist's claim about unaware experiences relies on. I will unpack these claims by first presenting some of the results of experiments on IB and CB and then explaining why the FOR theorist makes use of the same conception of experience that these studies undermine.

3.1 Change blindness

As was just noted, a common assumption about visual experience is that there is a richly detailed representation constructed for every moment or view. This representation is thought to capture the complete scene before one's eyes, including features that are and are not attended to by the perceiver and the metric properties of the entire scene. Within the constraints of the perceiver's perceptual systems, all the detail about the scene is present in the snapshot representation and the perceiver can choose to attend to certain details rather than others. On this view, whether the perceiver attends to any features present in the representation is irrelevant to the construction of the representation; attention is used to explore specific aspects of the wealth of detail present in the representation. Additionally, it is thought that a perceiver is able to detect changes in successive views of the same scene by comparing the detail of his current visual episode with the representation of his episode from just a moment before. Call the notion of visual experience involving a richly detailed, coherent representation of the entire visual scene the inner picture view and the idea that a detailed, metric representation of a scene survives across views of that scene the persistent picture view.

Recent studies of CB challenge the persistent picture view. Typically, change is perceived through the detection of associated “flickering” in the visual scene that attracts attention. However, if the perceiver is distracted by a global disturbance (or multiple extraneous transients) in the scene that occurs simultaneously with the local change, the local flickering is not detected and the perceiver does not notice the change. This effect can occur even when the local change happens in full view. It is not as though obscure elements of the scene go undetected in CB studies; large-scale changes of a sort that would otherwise be very easy to detect are completely missed by the perceiver. There are a number of CB demonstrations available on the World Wide Web that convincingly illustrate the phenomenon.¹³

As an example of the effects of CB, consider the following description of an experiment conducted by Kevin O’Regan, Ronald Rensink, and James Clark.¹⁴ Naïve subjects are instructed to watch a video display in which a sequence of images is shown. The first and last (i.e., original and modified) images are shown to the subjects for approximately 3 s, and the middle (i.e., distractor) image is displayed for 80 ms. The original image is of a car on the road, with a solid traffic line to its right. The distractor image is just like the original, but with “mudsplashes” (actually, small black-and-white patterned rectangles) occluding small parts of the image. Importantly, none of the mudsplashes occludes any part of the traffic line. The modified image is also the same as the first, except that the traffic line is now broken instead of solid. Although the mudsplashes create only a brief and rather minor overall disturbance, their effect on the perceiver’s ability to detect the change in the traffic lines is striking. Having done the on-line version of the experiment myself, it is clear that, without knowledge of what changes between the first and last images, one does not experience the change. Naïve subjects are not aware of the difference between the original and modified images. Once one knows what to look for,

however, the CB effect entirely disappears and the change becomes prominent.¹⁵ The test subjects were also presented with other, similar image sequences, and in some cases they noticed the change after seeing the images cycle through several times, whereas in other cases they never noticed the change during the 40 s interval that each image sequence was cycled.

CB experiments show that change cannot be consciously seen without attention. Such a result is incompatible with the persistent picture view of experience. In the absence of attention (which, in the example above was directed to a location of the scene other than where the change occurred), one does not notice the change, but once one knows what to look for and thus attends to the relevant location, one cannot help but see the change. If the persistent picture view were correct, the distractor should not have the effect of negating one's ability to detect change. The information from the immediately prior view of the same scene should still be accessible by the perceiving subject, and a comparison between the stored representation of the before state and the current visual experience of the after state should readily reveal what has changed. The fact that CB goes away immediately after one is told what changes in the scene highlights the role of attention in detecting change. Attention marks certain elements of the scene as salient and the rest are lost to consciousness. Those elements of the before view that have been indexed are loaded into visual short-term memory (VSTM), persist through the distractor event, and are available for comparison to the after view. If a richly detailed, coherent, metric-preserving representation is constructed for every visual episode, CB makes clear that very little of that information survives even a short while. The next phenomenon to be discussed, inattention blindness, makes it unlikely that such a representation is constructed in the first place, rendering the results of CB experiments not particularly surprising.

3.2 Inattention blindness

As change blindness experiments undermine the persistent picture view, research on IB has called into doubt the viability of the inner picture view itself. IB is the failure to detect the presence of visual stimuli that occur in plain view due to lack of attention. As is the case with CB, the features that go undetected in bouts of IB are significant and are such that, when IB subjects are confronted with what they did not notice during tests, they have no problem seeing what they missed and are shocked to learn that they did so. The result of IB research “is that there seems to be no conscious perception without attention” (Mack and Rock, 1998a, ix. Their emphasis). Although we may have the impression that our visual experience is richly detailed, containing far more information about the entire scene before us than we may ever be aware of, IB studies tell us that the only features we actually experience are those that we attend to. All other information about the visual stimulus is not present to consciousness.

Mack and Rock (1998a) conducted experiments in which subjects were engaged in tasks that demanded a high degree of attention, such as looking at a cross and trying to determine which arm is longer. After several trials of this task, an unexpected stimulus such as a small, black square was presented in one of the cross’s quadrants. Focused on their assigned task, a significant number of subjects did not notice the target object – approximately 25% of perceivers reported not noticing the unexpected object when it appears parafoveally and the cross is presented at fixation. Perhaps even more surprising is that when the cross is presented parafoveally and the unexpected stimulus appears at fixation, approximately 75% of the subjects reported not perceiving the target object. In subsequent trials, after being primed to look for the target object, almost every subject reported noticing the target object. Based on the evidence

from their tests, Mack and Rock argue that subjects did not perceive those items that they did not attend to and that attention is necessary for visual stimuli to become conscious.

Other research on attention and vision has turned up similar results. A significant feature of these studies is that they show that the failure to detect unexpected stimuli is attributable not only to attention being directed to another part (i.e., location) of the visual scene, but can also be attributed to focus of attention simply to other features in the same scene, including features that are visually co-located (i.e., occupy the same spot on the retina) with the unnoticed features. For example, Haines (1991) discusses a study of airplane landings using flight simulators equipped with head-up displays (HUD). The HUD is superimposed on the flight simulator's cockpit windshield and includes information usually presented in the plane's instrument panel. While focused on the HUD and their control of the aircraft during their landing approaches, two out of eight test subjects (who were commercial airline pilots) failed to notice that another airliner was sitting on the runway they were attempting to land on. The oblivious pilots blithely landed their planes "through" the parked airplane, completely unaware of what they had just done. In subsequent trials done without the HUD (i.e., using standard instrument panels), the pilots had no difficulty recognizing the parked plane. When informed of their virtual reality mishap during the HUD trials, the pilots were astonished and reported not having seen the large, stationary obstacle. This result seems to tell against Tye's earlier cited argument from his success at keeping his car on the road and his assumption that he was not paying attention to his visual sensations for the conclusion that he was subject to unattended visual experiences. The breakdown in attention paid to the visual scene that led the pilots to steer their aircraft directly into an obstacle would tend to support the revised reasoning about the distracted driver case, that

it follows from the fact that the driver is successful in keeping his car on the road that he was attending to his visual states.

Another relevant study with results consistent with the HUD study had subjects watch a 75 s videotape of a basketball-like game played in a small area, with the instruction that they should count the number of bounce passes and chest passes one of the two teams (one wearing black and the other wearing white) makes.¹⁶ In one test condition, in which separate video recordings of the two teams were superimposed in the videotape played for the test subjects, giving the players a “ghostly,” transparent appearance, 75% of the test subjects observing the black team and 92% of the test subjects observing the white team did not notice a person in a gorilla suit stroll across the court from left to right over a span of 5 s. Of particular interest is that at points during its path across the court, the person in the gorilla suit “passed through” the players and the ball, both of which the subjects were supposed to be attending to, yet went unnoticed. Due to the transparent quality of the scene displayed in the video, the person in the gorilla suit was occupying the same retinal location as the features of the scene that were being attended to, which means that the focus of the subjects’ attention had to be on something (or some feature) rather than a particular location in the visual field. In another test condition, filmed in live action with a single camera (giving the players a normal, opaque appearance) roughly half the test subjects failed to notice that a person in a gorilla suit stepped into the middle of the visual display, thumped its chest, and strolled off the court, all of which took 9 s and had the person in the gorilla suit both occluding and occluded by the players.

The results of these studies run counter to the idea that we experience more than we attend to and that our experiences are picture-like representations that can have certain features attended to while others go unnoticed. Instead, our experience represents only those features that

receive attention. From the standpoint of evolutionary engineering, this stands to good reason. Considering the limitations on processing power and bandwidth within our visual system, it would be a wise design choice to not create, much less store, richly detailed representations for features that are not of interest to the perceiver, with what is of interest being determined by what is attended to. Instead, our visual systems save their most valuable resources for recovery and processing of those features that best suit the perceiver's immediate conscious needs, while relying on assumptions about the relative stability of the external environment to obviate the need to create an exhaustive, resource-intensive representation for every view.¹⁷ In doing this, the visual system can also accommodate the unconscious acquisition of information about some features that have not been attended to, but which can unconsciously influence subsequent beliefs or actions, and which can draw the perceiver's attention, thereby bringing them into consciousness.

3.3 The inner picture view

Is the inner picture a mere straw man? Do the data from IB and CB studies actually have any negative consequences for the fortunes of FOR (or any other actual theory of experience)? As Kevin O'Regan notes, the inner picture view seems like something of an unfair caricature and does not appear to be explicitly endorsed by anyone.¹⁸ However, the FOR theorist is absolutely committed to such a view, if not explicitly, then certainly tacitly. FOR claims that perceptual states (of the right kind) are phenomenally conscious in and of themselves. A general statement of the representationalist thesis is that the phenomenal character of experience is a certain kind of representational content and that the relevant representational content includes features (i.e.,

properties, objects, and relations) of the stimulus scene. Thus the FOR theorist admits that our experiences – which are conscious level mental states - are structured entities containing a considerable amount of information about the stimulus scene and that we are free to attend to as much (within empirical limits) or as little of that information as we choose. Hence the analogy to pictures seems appropriate. As for whether anyone has come right out and endorsed the inner picture view, consider Tye’s claim that “the pertinent module [relevant to visual experience] is the one that has as its output a unified representation of the entire visual scene” (Tye 1995, 138).

I should note that I am not trying to convict the FOR theorist of embracing a theory of experience that involves a Cartesian Theater in which “inner pictures” are viewed. The FOR theorist, nor any other kind of representationalist, is not automatically committed to a privileged location in the brain in which experience “happens,” nor to a homuncular audience for experiences. My complaint is instead based on how the representations that are essential to phenomenal consciousness are to be understood. FOR theorists claim that the determination of the detail and structure of those representations is independent of what we are aware of in experience, whereas IB and CB indicate that which features and relations are represented is dependent upon what we attend to. It becomes increasingly difficult to take unattended features of the visual scene as phenomenally conscious especially in the light of the results of IB studies, as the IB subjects were incapable of reporting the presence of stimuli that occupy a significant portion of the visual scene and which at times were retinally co-located with stimuli that the subjects were phenomenally conscious of.

Perhaps one might say that ‘experience’ and ‘consciousness’ are being used here in a way different than the FOR theorist intends them to be used when it comes to our phenomenal awareness of things; such a response naturally suggests itself when one considers the remarks

FOR's proponents make on its behalf. However, a point to keep in mind is that in many important respects, inattentional blindness looks to be an attention-based counterpart to the scotoma-based phenomenon of blindsight. For example, like blindsighters, IB subjects cannot report the presence of certain stimuli in their visual field, yet further studies report that information about the blind area had been unconsciously acquired and can have an influence on the guidance of behavior or expectations of subsequent stimuli.¹⁹ If we take seriously the reports of blindsighters that they do not have visual experiences in their blind fields (which most, if not all, FOR theorists seem prepared to do) why do we not do the same for the IB test subjects regarding their reports that they do not see the unexpected stimuli they are presented? Although blindsight is a neural deficit condition and IB occurs as part of normal visual experience, there seem to be compelling reasons to conclude that one thing they share is that they result in a lack of phenomenal experience of certain features of a subject's visual field.

3.4 Unattended visual states lack poise

In spelling out his FOR theory of phenomenal consciousness, which he labels 'PANIC', Tye acknowledges that attentional deficits can prevent perceptual states from making an impact on the belief/desire system and he claims that attention must be "properly focused" on such states for them to have an impact (Tye 1995, 138). Even lacking attention, however, perceptual states of the right sort are claimed to still be phenomenally conscious; i.e., they are experiences with a distinctive phenomenology and which could have an impact on the belief/desire system, although they do not actually do so. I find the idea that unattended perceptual states nonetheless have felt qualities to be implausible, but I can see no way to articulate any reasons for my conviction that

go beyond mere reflection on the phenomenology of my own experiences. Fortunately, there are sound empirical reasons to doubt that “[states with PANIC] stand ready and available to make a direct difference to beliefs and desires” (Tye 2000, 172) in the absence of attention. In fact, the results of the CB and IB studies indicate that without attention directed to them, perceptual states are not poised (the ‘P’ in ‘PANIC’) in the way Tye requires them to be.

The empirical evidence shows that unattended features of the visual scene, even those that occupy the same retinal locations as attended features, are cognitively unavailable to us. As was noted in the previous section, they are consciously inaccessible by perceivers in much the same way that information about stimuli in their blind fields is unavailable to blindsighters. So, why should we count unattended features as even merely poised to impact (contrasted with actually impacting) our conscious-level deliberations? As demonstrated in the CB and IB studies, they are not introspectively available as elements of a richly detailed, scannable inner representation, since our visual experiences do not include such a representation. Furthermore, the opposite of what FOR predicts looks to be true – the construction of the representations relevant to experience is dependent upon where we direct our attention. Those features of our visual episodes that we do not attend to stand completely outside the realm of anything we should count as an experience, as it is not just that they do not affect conscious level deliberation, but that they are unable to. Unattended features can figure in the guidance of behavior only in an unconscious way, not taking a route through the belief/desire system, which is insufficient to qualify as being suitably poised. Given the close connection between experience and our rational faculties, there seems to be no reason to grant the status of experience to visual states (or their features) that go unattended.

Tye already has on offer a non-unitary FOR theory, in that his requirement of poise brings along with it the requirement that a perceiving creature must have higher-order states of some sort, if it is to qualify as a subject of phenomenal consciousness.²⁰ The PANIC theory looks promising, but it fails on empirical grounds if pitched as a FOR theory, even a non-unitary one. However, if the notion of poise is understood to include not only the requirement that the perceiving creature has higher-order cognitive states that can be impacted by its perceptual states, but also that for any perceptual state to be available to make an impact on beliefs and desires, it must have sufficient attentional resources allocated to it, the empirical difficulties raised in this note go away. By endorsing a form of HOR that draws on attentional resources to spell out the requirement of poise, the PANIC theorist seems to be giving up none of the advantages of the theory while avoiding one of its major problems.

4. Possible objections

Thus far, a considerable amount of data has been presented in an effort to show that FOR has empirical commitments it cannot meet and that an adequate HOR theory should be grounded in attentional resources. The view developed in this paper is susceptible to challenges from both FOR theorists generally and from competing positions within the HOR camp. At this point, I would like to address two avenues of attack the FOR theorist might pursue against HOA, while leaving internecine disputes with other HOR theorists to the side for now. The objections entertained below might, in some similar form, also be raised by a HOR theorist, but it seems to me far more likely that they would be lodged by a FOR theorist.

4.1 Do the empirical data support such a strong conclusion?

Since the case made against FOR in this note is grounded in empirical data, an obvious starting point for any FOR theorist's rebuttal would be to question how much support the cited data actually provide for the anti-FOR conclusion. In particular, given the nature of the studies referred to, one might complain that the most that has been shown is that we sometimes are phenomenally unconscious of unattended features of the visual scene, not that we always are phenomenally unconscious of unattended features of the visual scene. One seems forced to make a leap from, say, the results of Mack and Rock's IB experiments to the claim that there is no conscious perception in the absence of attention. Perhaps the FOR theorist could successfully make the argument that there is no good reason to make that leap and thus that the results of research on the role of attention in conscious experience do not militate on behalf of HOR.

This objection can be sharpened by noting that we sometimes, perhaps even often, seem to observe more than we attend to. For example, if I look out the window of my office at the city skyline and then close my eyes, it would seem that I can recall visual details about the skyline that I had not attended to. For another example, I might suddenly notice that my doorbell is ringing and that, in fact, it has already rung several times before I noticed (i.e., attended to) the ringing. Additionally, sudden changes in the visual scene can draw our attention to previously unattended features. Since the detection of change occurs through the detection of "flickering" in the visual scene, presumably that detection of flickering is conscious (since we are conscious of the change) yet occurs without being attended to.

An important point to keep in mind when considering such examples is that, according to the IB paradigm, perceptual activity that occurs in the absence of attention is not conscious, but from that it does not follow that unattended perceptual activity does not result in information

pickup, or that perceptual activity outside the realm of attention is unstructured or amorphous. This runs squarely opposite to the conclusions drawn by several prominent researchers, who contend that attention is required to bring structure to or “bind” features of the visual scene.²¹ Mack and Rock posit a distinction between conscious and unconscious perception, and claim that attended and unattended stimuli are both processed the same way, including the construction of an implicit representation of the external world. When a representation reaches a certain threshold of activation, it captures attention and becomes conscious. Thus perceptual sensitivity and activity are not one and the same as perception. As Noë and O'Regan (2000) put it, “[to] experience detail, one must detect it. But to detect it, there is no requirement that one experience it.”

Empirical studies were cited earlier in which unattended information that is acquired through perceptual activity was used to guide behavior. Additionally, as is also the case in blindsight, implicit representations that never reach the threshold of activation required for consciousness nonetheless can have an effect on subsequent conscious-level activity such as decisions and behavior.²² It is plausible to think that examples of the sort presented above can be handled within the IB and CB paradigms by claiming that the representation constructed on the basis of the subject’s perceptual activity was initially unconscious (because it was not attended to) and it either (a) eventually reached a critical threshold and attracted attention, thereby becoming conscious, or (b) it never reached the critical threshold but was still in a position to influence the subject’s conscious-level cognitive and behavioral activity. (a) would seem sufficient to cover both the cases of the sudden realization that the doorbell has been ringing for some time and the detection of change in previously unattended features of the visual scene. The doorbell’s continued ringing eventually causes the subject’s previously unconscious auditory

representation to reach a level of activity that (finally) attracts attention. The detection of flickering in the visual scene used to detect change does so by drawing attentional and processing resources to transients in the visual scene, which brings the associated change to consciousness. (b) applies to cases in which features we did not attend to can affect our subsequent impression of what was previously before our eyes, such as my apparent recollection of unattended features of the city skyline. The IB experiments show that our impression that we consciously take in more of the visual scene than we attend to is false, but there is still an opportunity for unattended – and unconscious – information about the visual scene to impact our conscious level reflection on the visual scene, thus leading to the sense that one experiences more detail of the visual scene than is attended to.²³

Note that in the case of change detection, the detection of the associated flickering need not itself be conscious – in fact, it is an automatic mechanism of our visual systems²⁴ – but the result of that detection, the drawing of attention to the relevant features of the visual scene, is what makes us conscious of the change. Returning to the quotation from Noë and O'Regan above, to experience change we must detect it, but we need not experience change to detect it. In the CB experiments, the subjects failed to notice the large-scale changes in the scene before them because they were not attending to the relevant features of the scene and they (actually, their visual systems) could not detect the associated flickering of the unattended features due to the overwhelming nature of the simultaneous flickering associated with the distracting event. The subjects' flicker detection mechanisms - and thus their attentional and processing resources - were drawn not to the local transients, but instead to the global transients or were interfered with altogether. Thus the change blindness subjects failed to experience the local (yet significant)

changes not because they did not experience the associated flickering, but rather because their attention was automatically directed to some other feature(s) of the visual scene.

Moving away now from the examples and returning to the more general challenge to the sufficiency of the empirical data, it does seem fair to ask whether the IB and CB experiments establish that there is no experience in the absence of attention. Rensink (2000), for one, has granted that the CB experiments are successful in establishing that attention is needed to experience change, but he thinks that the IB experiments fall short of securing that experience of features of the visual scene requires attention. Rensink's concern about what the IB research establishes is not based on counter-examples or conflicting evidence, but rather limitations on devising test procedures that would establish both that when a stimulus is experienced, attention is present and that when a stimulus is not experienced, attention is missing. Further experimental procedures are needed to validate the conclusions Mack and Rock draw on the basis of their evidence. I think Rensink's worry is well placed. Nonetheless, I do think that the results of the IB experiments are quite striking and lend a great deal of support to the strong (and admittedly controversial) claim that experience requires attention. It is entirely possible that this is an issue that cannot be settled decisively through any set of experiments we might devise and that we will simply have to rely on the marshaling of a considerable amount of empirical data that all converge on the same point to guide us in the conclusions we draw about the relation between attention and experience. However, there is no reason at this point to think that all possible experimental methods that would confirm Mack and Rock's claim have been exhausted and we certainly have no good reason on hand for rejecting it altogether.

Considering that the IB research does demonstrate that in (at least) some circumstances, subjects do not experience significant features of the visual scene in the absence of attention,

even if the necessity of attention for all experience cannot be decisively established, researchers and theorists working on the nature of experience must take into account the results of the IB experiments. Suppose for a moment that it were to turn out that there are situations in which subjects have experience in the absence of attention. We would still have to explain the difference between those situations and the situations in which subjects do not experience what they do not attend to. Perhaps we would have to countenance different kinds of experience, some that are attention-based and others that are independent of attention. Obviously, this sort of speculation is of limited value. What is clear from the available evidence, however, is that a FOR-only theory (i.e., one that covers all the kinds of experiences we might uncover) is untenable, while some version of HOR grounded in attention will turn out to be correct. Thus the HOA theory proposed in this note will either be the whole story about experience or a major part of an overarching theory that blends together HOR about certain kinds of experience and FOR about other kinds.

4.2 Is HOA really a higher-order theory?

The empirical data referred to in this note show that a particular way of characterizing FOR, that perceptual states with an appropriate pedigree are intrinsically conscious, is false.²⁵ Without attentional resources attached to them, perceptual states are not conscious. However, a FOR theorist might respond that such a characterization is incomplete and that the more central issue is whether perceptual states are conscious only as the objects of higher-order states. Perhaps the FOR theorist could admit that conscious perceptual states require attentional resources attached to them, but that the relevant sort of attention is of a first-order sort, not a higher-order sort.

Roughly, the idea would be that attention must be directed to features of the visual scene for them to be conscious, but one does not have to (and perhaps cannot) direct attention to experiences themselves.

The most promising way the FOR theorist might develop this response is to mark a distinction between experience and the content of experience. So, in the case of visual experience, in order for a visual state to be conscious, the perceiver is not required to attend to the visual state itself, but rather to the external features of the visual scene represented by the visual state. Such attention might reasonably be thought to be first-order, since it does not seem to consist of one representational state having another as its object. To take one FOR theorist as an example, we have seen that Tye is on record as claiming that the distracted driver is “not aware of his visual sensations. He is not paying any attention to them” (Tye 1995, 115). I read this and other comments Tye makes about (not) attending to visual sensations to mean that the subject is not attending to his visual states, but they are nonetheless phenomenally conscious.²⁶ Tye could continue to take such a stance while modifying his account so that what is relevant to whether the driver’s visual states are phenomenally conscious is whether he attends to features included in the contents of those states. He could do so while maintaining his FOR position by saying that the sort of attention required to make the driver’s visual states conscious is directed not at the visual states themselves, but rather at the features of the world that the experience is an experience of. Thus a perceptual state becomes conscious not as the object of a higher-order representation, but through having elements of its representational content attended to.

A response to this revision of FOR that immediately comes to mind is that it seems quite odd to suppose that one could attend to a feature represented in one’s visual state without attending to one’s visual state itself. It is not as though attention has direct access to things in the

environment; attention to things in the environment is mediated by states of our perceptual systems. The things in the world attended to are available as objects of attention only because they figure in the contents of perceptual states. Thus to attend to some feature F of visual state S, S itself would have to be attended to, for how else would I have access to F? There is some merit to this sort of response, but it gains more bite when paired with certain observations about how the representationalist position should be developed, particularly in regard to the issue of the transparency of experience.

The proposed revision to FOR relies on a distinction between experiences and their contents, so that it is possible for a perceiver to attend to features of the visual scene without attending to her visual state itself. However, it is questionable to think a representationalist can make a legitimate appeal to such a difference. Consider the appeal to transparency often cited by representationalists.

[Introspection] of your perceptual experiences seems to reveal only aspects of what you experience, further aspects of the scenes as represented. Why? The answer, I suggest, is that your experiences have no introspectible features over and above those implicated in their intentional contents. (Tye 1995, 137)

One of the main things that sets representationalism apart from qualia freakery is that representationalism has it that the individuation of experiences does not depend on consciously accessible properties of experiences themselves; there are no such intrinsic properties of experience to be found when a perceiver turns her “inner eye” on her experiences themselves. Experiences are individuated solely on the basis of their representational contents. Why, then, should we mark any sort of useful distinction between experiences and their contents? As far as the representationalist is concerned, there is no detectable difference between an experience and

its content – in being aware of an experience, one is not aware of the experience itself as an object, but rather the objects, properties, and relations that it is an experience of. That does not entail, however, that one is incapable of attending to one’s visual states, only that when one attends to one’s visual state, all that one will find are the features that figure in the state’s content.

The passage from Tye quoted above is consistent with such thinking, in that he seems to be saying that one can introspect one’s experiences, but all that one will find when one does so are the features of the content of the experience introspected. Furthermore, given the apparent collapse of a distinction between experiences and their contents, when attending to elements of an experience’s representational content, you are thereby also attending to the experience itself, since that is all that is cognitively available to the subject. Applying this to the current proposal regarding attention, we have the appearance of first-order attention when what is really taking place is a form of higher-order attention.

In questioning whether the relevant attentional states really are higher-order, the advocate of the proposed revised version of FOR would also seem to be calling in to question whether we are really able to have higher-order thoughts about our experiences. One’s thoughts about one’s experiences also always “slip through” the experiences themselves and on to the features represented in the experience; in trying to form a higher-order thought about one’s experiences, what one ends up with are thoughts about the things in the world that one’s experience is an experience of. If one disputes this and contends that we are able to form higher-order thoughts about experiences themselves (separate from what they represent), our ability to do so could only be explainable by the presence of detectable non-representational properties of experience, which runs directly counter to the representationalist’s position. Even though HOA grants that higher-

order thought about perceptual states is not required for experience, it certainly seems plausible to suppose that we are capable of forming higher-order thoughts about our experiences. Even a committed FOR theorist such as Tye thinks that we routinely form such thoughts.²⁷ If one is willing to count as higher-order those thoughts about one's experience that turn out to be thoughts about things in the world, then why should the same not also apply to the attentional states that are directed towards one's experiences that turn out to be directed towards things in the world? Anyone interested in attempting to rescue FOR by challenging whether the attentional states implicated by HOA really are higher-order is faced with one of two daunting tasks: (1) make a plausible argument that we cannot form higher-order thoughts about our experiences, thereby cutting off the appeal to the possibility of higher-order thought about experience to justify taking the relevant attentional states to be higher-order; or (2) admit that we can have higher-order thoughts about experience and explain how the way in which those thoughts get counted as higher-order cannot be extended to the attentional states required for experience.

5. Conclusion

FOR's claim that there can be experience in the absence of higher-order representation of one's sensory states has testable, empirical commitments. The data from studies on the cognitive demands of driving, change blindness, and inattention blindness indicate that FOR is not an empirically plausible theory of consciousness. However, the HOR theory we are led to by consideration of the empirical facts about experience is not one that demands that the relevant higher-order states are thought-like or perception-like. Instead, one promising way to pursue

filling out the details of the proposed HOA theory is to take an existing FOR theory, Michael Tye's PANIC theory, and modify it so that the required connection between our perceptual and rational faculties depends, at least in part, on attentional resources being dedicated to the outputs of our perceptual systems. Although it remains a further empirical question just how far down the phylogenetic ladder the relevant attentional mechanisms extend, it is prima facie plausible to suppose that they extend far enough to at least capture the central examples of the kinds of creatures that a theory of consciousness should not rule out as subjects of experiences; e.g., human infants, other primates, dogs, and cats. Thus HOA looks well-positioned to be developed in a way that accommodates both the empirical facts about the cognitive demands of experience and our intuitions about what sort of creatures should be counted as subjects of experiences.

6. Notes

¹ The operative, and quite plausible, assumption throughout this discussion being that a necessary condition of a state being an experience is that it is conscious.

² For FOR theories, see Tye 1995 and 2000, Dretske 1995, and Kirk 1994. For HOT theories, see Rosenthal 1986 and 1993, and Carruthers 1998. For HOP theories, see Armstrong 1980 and Lycan 1987 and 1990.

³ One important difference between HOA and HOP is that on the former, but not that latter, there is no mystery why the relevant higher-order states have a "transparent" phenomenology. Attentional states are not the kinds of things one would expect to possess a phenomenal character of their own, whereas it is puzzling that states of an inner sense (or "scanner") would not have their own distinctive phenomenology rather than the phenomenology of the perceptual states

they detect; see Dretske 1995, 62. There are obvious parallels between the position I take on consciousness and attention and William Lycan's HOP view. I will not discuss in this paper the points on which I believe Lycan and I differ, and it should suffice for present purposes to group our views together, if that is helpful to the reader. I should note, however, that I do not read Lycan to be a true HOP theorist, despite his claim to be one, but as having a view closer to the HOA theory discussed here, as his way of explaining how the HOP theorist avoids being committed to a Cartesian Theater account of the mind appears to distance him from a commitment to an internal perception-like mechanism (see Lycan 1997, 761-3).

⁴ It is important to point out that Armstrong is a HOP theorist, not a FOR theorist. I mention his presentation of the distracted driver case while criticizing the FOR theorist's handling of it to support my contention that HOR theorists and FOR theorists traditionally have not differed in how they pitch the distracted driver scenario, but rather in how they interpret it.

⁵ Tye 1995, 138 and 2000, 62. The notion of poise Tye has in mind has to do with standing ready to make an impact on the belief/desire system. A poised state need not actually make an impact. In section 3, I will take up Tye's claim that a state can be poised despite "attentional deficits ... preclud[ing] belief formation" (Tye 1995, 138).

⁶ See Treisman 1998.

⁷ See Noë and O'Regan 2000. It is likely, of course, that the conscious-level operation of these attentional mechanisms also involves the activity of subconscious attentional mechanisms.

⁸ I should point out that as an account of what makes a sensory state a candidate for consciousness (i.e., the kind of sensory state that, when attention is directed to it, is conscious), I

would endorse something along the lines of Tye's PANIC theory. I shall have more to say about the PANIC theory at the close of the next section.

⁹ See Bridgeman et al 1997, Aglioti et al 1995. A clear statement of how unconscious information can guide behavior despite being in conflict with what is consciously experienced is found in Aglioti et al 1995, 684:

... the calibration of grip aperture is quite refractory to the compelling size-contrast illusion induced by the [experimental] display. This result suggests that the automatic and metrically accurate calibrations required for skilled activity may depend on visual computations that are different from those driving our perceptual judgments about objects in the world. In short, what we think we see may not always be what guides our actions.

¹⁰ For more on change blindness, see O'Regan 1992, Rensink 2000, Simons 1996. For more on inattention blindness, see Mack and Rock 1998a and 1998b, O'Regan and Noë 2001, Simons 1999 and 2000, Simons and Chabris 1999. Several of these references discuss both topics.

¹¹ Simons and Chabris 1999, 1060.

¹² Kevin O'Regan refers to this conception of experience as the "internal screen"; see O'Regan 1992.

¹³ The reader is encouraged to go to Daniel Simons' Change Detection Database at <http://viscog.beckman.uiuc.edu/change/demolinks.shtml> and Kevin O'Regan's webpage at <http://nivea.psychology.univ-paris5.fr/index.html> to view some of the available demos.

¹⁴ The results of the experiment are discussed in O'Regan et al 1999. The reader can see this demonstration in action at http://nivea.psychology.univ-paris5.fr/Mudsplash/Nature_Supp_Inf/Movies/traf.GIF.

¹⁵ This is not always the case. See Rensink et al 1997 for discussion of CB that persists even when subjects expect change.

¹⁶ See Simons and Chabris 1999.

¹⁷ See O'Regan 1992 for more on the external world as a memory store.

¹⁸ See O'Regan 1992.

¹⁹ See Simons 1999, 166; Mack and Rock 1998b, 72-3.

²⁰ See Lurz 2000.

²¹ Perhaps most prominent among the researchers advocating such a model is Anne Treisman. See Treisman 1998 and Treisman and Gelade, 1980.

²² See Simons 1999, 168.

²³ O'Regan and Noë discuss in several places other factors, especially our visual system's assumptions about the relative stability of our environment, that give rise to the sense that we experience a good deal more than we attend to.

²⁴ See O'Regan and Noë, 2001

²⁵ Given the discussion of section 4.1, of course, one might wish to read this claim as restricted to the types of experience treated in IB and CB experiments, leaving open whether the claim applies to experience in general.

²⁶ For another example of Tye's usage of 'visual sensation' in this way, see Tye 2000, 14.

²⁷ See Tye 1995, 115.

7. References

- Aglioti, S., DeSouza, J. F., & Goodale, M. A. 1995. Size-contrast illusions deceive the eye but not the hand. Current Biology, 5, 679-685.
- Anderson, D., Abdalla, A., Pomietto, B., Goldberg, C., & Clement, C. 2002. Distracted driving: review of current needs, efforts, and recommended strategies. Senate Document no. 14, Commonwealth of Virginia.
- Armstrong, D. 1968. A Materialist Theory of Mind. London: Routledge.
- Armstrong, D. 1980. The Nature of Mind and Other Essays. Ithaca, NY: Cornell University Press.
- Bridgeman, B., Peery, S., & Anand, S. 1997. Interaction of cognitive and sensorimotor maps of visual space. Perception & Psychophysics, 59, 456-469.
- Carruthers, P. 1998. Natural theories of consciousness. European Journal of Philosophy, 6, 203-222.
- Dretske, F. 1995. Naturalizing the Mind. Cambridge, MA: MIT Press.
- Haines, R. 1991. A breakdown in simultaneous information processing. In L. Stark & G. Obrecht (Eds.), 4th international Symposium on Presbyopia, 171-176. New York: Plenum.
- Harbluk, J.L., Noy, Y.I., & Eizenman, M. 2002. The impact of cognitive distraction on driver visual behavior and vehicle control. TP# 13889 E, Transport Canada.
- Kirk, R. 1994. Raw Feeling. Oxford: Clarendon.

- Lee, J., Caven, B., Haake, S., & Brown, T. 2000. Speech-based interaction with in-vehicle computers: the effect of speech-based e-mail on drivers' attention to the roadway. Technical Paper 27, National Highway Transportation Safety Administration.
- Lurz, R. 2000. A defense of first-order representationalist theories of mental-state consciousness. Psyche, 6.
- Lycan, W. 1987. Consciousness. Cambridge, MA: MIT Press.
- Lycan, W. 1997. Consciousness as internal monitoring. In N. Block, O. Flanagan, & G. Guzeldere (Eds.), The Nature of Consciousness, 755-771. Cambridge, MA: MIT Press.
- Mack, A. & Rock, I. 1998a. Inattentional Blindness. Cambridge, MA: MIT Press.
- Mack, A. & Rock, I. 1998b. Inattentional blindness: perception without attention. In R.D. Wright (Ed.), Visual Attention, 55-76. New York: Oxford University Press.
- Noë, A. 2002. Is the visual world a grand illusion? Journal of Consciousness Studies, 9, 1-12.
- Noë, A., & O'Regan, J. K. 2000. Perception, attention, and the grand illusion. Psyche, 6.
- O'Regan, J. K. 1992. Solving the 'real' mysteries of visual perception: the world as an outside memory. Canadian Journal of Psychology, 46, 461-488.
- O'Regan, J. K. & Noë, A. 2001. A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences, 24, 955-975.
- O'Regan, J. K., Rensink, R. A., & Clark, J. J. 1999. Change-blindness as a result of 'mudsplashes'. Nature, 398, 34.
- Rensink, R. 2000. When good perceivers go bad: change blindness, inattentional blindness, and visual experience. Psyche, 6.

- Rensink, R. A., O'Regan, J. K. & Clark, J. J. 1997. To see or not to see: the need for attention to perceive changes in scenes. Psychological Science, 8, 368-373.
- Richardson, D.C. & Spivey, M.J. 2000. Representation, space, and Hollywood Squares: looking at things that aren't there anymore. Cognition, 76, 269-295.
- Rosenthal, D. 1986. Two concepts of consciousness. Philosophical Studies, 94, 329-359.
- Rosenthal, D. 1993. Thinking that one thinks. In M. Davies and G. Humphreys (Eds.), Consciousness, 207-211. Oxford: Blackwell.
- Simons, D.J. 1996. In sight, out of mind: when object representations fail. Psychological Science, 7, 301-305.
- Simons, D.J. 1999. To see but not to see. Journal of Mathematical Psychology, 43, 165-171.
- Simons, D.J. 2000. Attentional capture and inattention blindness. Trends in Cognitive Sciences, 4, 147-155
- Simons, D. J. & Chabris, C. F. 1999. Gorillas in our midst: sustained inattention blindness for dynamic events. Perception, 28, 1059-1074.
- Treisman, A. 1998. The perception of features and objects. In R.D. Wright (Ed.), Visual Attention, 26-54. New York: Oxford University Press.
- Treisman, A. & Gelade, G. 1980. A feature-integration theory of attention. Cognitive Psychology, 12, 97-136.
- Tye, M. 1995. Ten Problems of Consciousness. Cambridge, MA: MIT Press.
- Tye, M. 2000. Consciousness, Color, and Content. Cambridge, MA: MIT Press.