

Attention, Visual Consciousness, and Indeterminacy

James Stazicker

(Revised March 2010)

Abstract: I propose a new argument showing that conscious vision sometimes depends constitutively on conscious attention. I criticise traditional arguments for this constitutive connection, on the basis that they fail adequately to dissociate evidence about visual consciousness from evidence about attention. On the same basis, I criticise Ned Block's recent counterargument that conscious vision is independent of one sort of attention ('cognitive access'). Block appears to achieve the dissociation only because he underestimates the indeterminacy of visual consciousness. I then appeal to empirical work on the interaction between visual indeterminacy and attention, to argue for the constitutive connection.

1. Conscious Attention and Conscious Vision: How are they Related?

William James defined attention by appeal to experience and folk psychology:

Every one knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others.

(James 1890/1950, p.403)

To attend to something, in this folk-psychological sense, is to be conscious of it in a certain way. To see something is also to be conscious of it in a certain way – to be visually conscious of it, we might say. What's more, you often attend to the very things you see. For example, you're currently not only seeing the words on this page but also attending to them, focussing your consciousness on them in order to read.

For helpful discussions and comments, thanks to Ned Block, Michael Caie, John Campbell, Zoe Drayson, Geoffrey Lee, Mike Martin, Alva Noë, Ariel Rokem, John Schwenkler, Michael Silver and members of his lab, and an anonymous referee.

Address for correspondence: Department of Philosophy, Moses Hall, U.C. Berkeley, Berkeley CA 94720, U.S.A.

Email: stazicker@berkeley.edu

What, then, is the relationship between attention and visual consciousness? It's clear that attending causes episodes of conscious seeing. For example, if something takes your interest you can direct your gaze at it, and thereby come to see its details. But is there also a closer connection? In particular, is attention constitutive of visual consciousness?

Psychologists distinguish between overt attention, which in the visual case involves movements of the head or eye, and covert attention, which involves no such observable behaviour. I'll explore the question whether *covert* conscious attention is constitutive of visual consciousness. We could also ask a parallel question about overt conscious attention. But that question would concern, *inter alia*, the relationship between conscious vision and a certain sort of bodily behaviour, and that is not my concern here. Covert attention is often described as operating like a spotlight, concentrating awareness and information-processing where its beam falls (e.g. Posner, 1980). We might spell out this analogy in either of the following ways:

Independence

Attention ranges like a spotlight over the objects you see in the environment around you, concentrating consciousness on them. In that sense, attention illuminates objects whose character is independent of it. Likewise, attention ranges over the contents of visual consciousness, illuminating episodes of conscious vision without constituting them.

Dependence

Attention ranges like a spotlight over the objects you see in the environment around you, concentrating consciousness on them – in that sense illuminating them. But visual consciousness is not always fixed independently of attention, there in advance

to be illuminated by it, so to speak. Rather, attention is sometimes one of the factors constituting an episode of conscious vision and the way it presents things.

Of course, we could define a notion of visual experience such that *Dependence* is true. We could choose to treat visual experience as an episode which includes the conscious effects of attention. But that would avoid the interesting question about the structure of experience. It would invite the response that there is a further conscious visual episode, constituted independently of attention, which is only one element in visual experience as we're understanding it. A more promising approach is to tackle the question head-on, asking whether there is in fact always some further conscious visual episode constituted independently of attention.

What is it for one conscious episode to constitute another, or for one episode to depend constitutively on another? Sometimes philosophers use the term 'constitutive' to mean *essential*. For example it's argued that being of chemical composition H₂O is constitutive of, or essential to, being water. But the term is also used to express a weaker connection – weaker in the sense that it's entailed by but doesn't entail essential connection. Scientists and philosophers investigating the neural basis of consciousness sometimes use 'constitutive' to describe a *minimal sufficient condition*: the constitutive basis of a particular conscious episode is the set of factors jointly sufficient for the phenomenal character and representational content of the conscious episode, and such that no proper subset of this set of factors is so sufficient. For example, following Cristof Koch, Ned Block tries to identify the 'minimal set of neuronal events and mechanisms jointly sufficient for a specific conscious percept' (Koch, 2004, p.16; Block, 2005, p.46). Block equates this set of events and mechanisms with the 'minimal neural basis of the phenomenal *content* of an experience, that which differs between the experience as of red and the experience as of green'.

This constitutive connection between physical and conscious phenomena may also entail constitutive connections *among* conscious episodes: if the physical basis minimally sufficient for a conscious episode *A* includes the physical basis minimally sufficient for a conscious episode *B*, then *B* is itself part of the minimal sufficient condition on *A*; *B* is among the set of factors sufficient for the character and content of *A*, and such that no proper subset of this set of factors is so sufficient; *A* depends constitutively on *B*. I'll pursue the question whether, in this sense, visual consciousness sometimes depends constitutively on covert conscious attention.

To see the sense in which a minimal sufficient condition is a constitutive condition, consider a mountain that's partly made of – or partly constituted by – some granite. Being made of granite is not essential to being a mountain. Plenty of mountains have no granite in them. Nevertheless, the granite is constitutive of this particular mountain in an interesting sense:¹ the granite is among a set of factors jointly sufficient for the appearance, structure and causal powers of the mountain, and no proper subset of this set of factors is so sufficient. This same notion can be applied to events and episodes as well as to objects. The 1906 eruption of Vesuvius was constituted in part by an expulsion of magma from some rock, in that the expulsion of magma was among a set of factors jointly sufficient for the appearance, structure and causal powers of the eruption, and such that no proper subset of this set of factors was so sufficient.²

¹ Or more than one interesting sense: the granite is a spatial part of the mountain; this connection may entail the further, more abstract constitutive connection I elaborate. I focus on the more abstract connection, because it applies more obviously to conscious episodes.

² We should accept this constitutive connection whether or not we also accept, with Davidson, 1969, that the 1906 eruption of Vesuvius is *identical* with a complex event of which the expulsion of magma is a part.

It's controversial exactly how sufficiency should be understood here. Konrad Marc-Wogau, 1962, introduced the notion of a minimal sufficient condition into the philosophical literature to capture the historian's idea of a cause. Koch and Block have in mind a modally stronger condition, a condition on a conscious episode which can be contrasted with its mere causes. The simplest proposal would be that *A* is minimally sufficient for *B* only if it's absolutely impossible for *A* to occur without *B* occurring. But it's controversial whether physical factors are sufficient in this way for conscious episodes. Perhaps minimal sufficiency can be understood in a modally weaker way which nonetheless preserves the distinction between constitutive or minimal sufficient conditions on the one hand, and mere causes on the other. Specifying such a condition would be a substantial project, and I don't want to get involved in that here. I'll take *Dependence* to claim, and *Independence* to deny, that conscious attention is part of a minimal sufficient condition on some episodes of conscious vision, where the sufficiency here is modally as strong as the sufficiency connecting physical factors with conscious episodes – however strong exactly that turns out to be.

I think progress here requires that we give centre-stage to the indeterminacy of visual consciousness. In what follows, I'll show how underestimating visual indeterminacy lends spurious support to *Independence*, and I'll show how taking visual indeterminacy seriously makes it plausible that *Dependence* is correct.

Independence is perennially tempting, and there is a danger that it gets presupposed in our philosophical theories about visual experience. Consider Mike Martin's analysis:

The notion of a visual experience of a white picket fence is that of a situation being indiscriminable through reflection from a veridical visual perception of a white picket fence as what it is.

(Martin, 2006, p.363)

Martin intends this style of analysis to apply to visual experiences in general: to have a visual experience as of a phenomenon φ is to be in a situation which is indiscriminable through reflection from a veridical visual perception of φ as what it is. 'Reflection' consists in attending to your experience, and so in attending to the objects of that experience (cf. Martin, 1998). So on the face of it at least, the analysis presupposes *Independence*. To see this, suppose that *Dependence* is correct: when you attend to something you experience φ , and when you don't attend to it you experience χ . Martin's analysis seems to rule out this scenario, treating both the experience of φ and the experience of χ as experiences of φ . For reflection on your visual experience of χ will alter it into a visual experience of φ , hence into an experience indiscriminable from a veridical visual perception of φ as what it is.

Merleau-Ponty railed against *Independence*. Here's his critical characterisation of it:

The function which reveals [the contents of conscious vision], as a searchlight shows up objects pre-existing in the darkness, is called attention. ... Inattentive perception contains nothing more than and indeed nothing other than the attentive kind. ... We are not called upon to analyse the act of attention as a passage from indistinctness to clarity, because the indistinctness is not there.

(Merleau-Ponty, 1945/1962, pp.30-32)

According to Merleau-Ponty, attending in fact suppresses the 'indeterminate ... quality' of inattentive vision. I doubt that *Dependence* could be established just through reflection on our experiences, as Merleau-Ponty suggests. I'll offer an empirical argument instead. But my argument has an important point of contact with Merleau-Ponty: I'll argue that attention is sometimes constitutive of an increased determinacy in visual consciousness.³

³ As an anonymous referee notes, further points of difference are that Merleau-Ponty is concerned with overt attention as much as with covert attention, and that his notion of indeterminacy differs from the one I'll define.

I won't offer any necessary conditions on visual consciousness, but I will work with the following sufficient condition. Sometimes we think about what we see in a distinctive, non-descriptive way, which we can express by talking about '*that* object' or '*that* shape', for example. This kind of visual-demonstrative thought is made possible by visual consciousness, and is possible only where there is visual consciousness of the phenomenon selected in thought (Campbell, 2002). We can use the same forms of words to express thoughts which do not exploit visual consciousness in this way, but those thoughts are *ipso facto* not visual-demonstrative thoughts. Although this is in one sense a trivial sufficient condition on visual consciousness, we'll see that it's very useful in practice. It's important that I'm not assuming that visual-demonstrative thought is necessary for visual consciousness. Thoughts of this kind involve attention to what you see, and it's an empirical question whether conscious seeing requires attention.⁴

Similarly, I'll take the condition James describes to be sufficient for attention, without assuming that it's necessary. In the contemporary cognitive sciences, the term 'attention' is used to refer to various sorts of selective information-processing. From this perspective, it's controversial whether attention always involves conscious awareness (Koch and Tsuchiya, 2006). So James stipulates the outcome of what looks like a substantive debate, by defining attention in terms of a conscious 'essence'. Still, we're all familiar, as James says, with attention as a conscious phenomenon – whether or not the same phenomenon also occurs non-consciously. And much of the scientific research into attention does proceed by manipulating this conscious phenomenon. So if we're careful we can draw on that research to learn about conscious, Jamesian attention.

⁴ To engage with the work I criticise, I'll treat visual consciousness as a form of representation. My arguments could be re-cast in different frameworks – e.g. Campbell's (2002) account of visual consciousness as a relation to objects and properties.

Visual attention can be controlled *exogenously*, i.e. directed by automatic effects of stimuli on attention, or it can be controlled *endogenously*, i.e. directed by your aims and decisions. Some theorists argue that these two sorts of attention do not form one natural kind. For example, Prinzmetal and Landau, 2008, claim that the effects and mechanisms of endogenously controlled attention differ from those of exogenously controlled attention. I'll return to this issue in Part 4. For now, note that the question whether *Independence* or *Dependence* is true does not turn on it: if there are two distinct kinds of conscious attention, then *Independence* says that visual consciousness is constitutively independent of both them, and *Dependence* says that visual consciousness depends constitutively on at least one of them.

In Part 2, I'll explain some special problems which the relationship between visual consciousness and attention poses for experimental investigation, and criticise a popular line of argument for *Dependence*. In Part 3, I'll criticise Ned Block's most recent argument that visual consciousness is independent of the sort of attention involved in verbal report, on the basis that Block underestimates the indeterminacy of visual consciousness. In Part 4, I'll appeal to attention's role in fixing the determinacy of visual content, to argue that attention is sometimes constitutive of visual consciousness.

2. Evidence for Visual Consciousness Involves Attention

Arien Mack and Irvin Rock, 1998, argue for an extreme version of *Dependence*: the striking claim that 'attention is essential to conscious perception'; we do not consciously experience

unattended objects of perceptual processing. They call this phenomenon 'inattention blindness'.⁵

Mack and Rock's experimental method is illustrated in Figure 1. Subjects were given a task to occupy their attention. They had to report the longer arm of a cross. (Note that what's manipulated here is exactly the familiar sort of conscious attention involved in selecting something you see for thought.) The cross was presented for 0.2 seconds either at fixation or peripherally within 2.3° of fixation.⁶ Fixation was held constant while covert attention was manipulated in this way. This process was repeated in four trials, except that in either the third or the fourth trial a further stimulus accompanied the cross, without warning. This stimulus was located in a quadrant of the cross, either at fixation or within 2.3° of fixation, depending on the position of the cross. On this critical trial, once subjects had reported the longer arm of the cross, they were immediately asked also to report whether they had seen anything that had not been present on previous trials, and if so to identify it or select it from an array of four to six objects. The idea was that these latter reports would reveal what subjects were visually conscious of, under a condition of inattention. Various stimuli were used, to test visual consciousness of various phenomena under a condition of inattention. These stimuli were all clearly visible when subjects attended to them: in a full-attention control trial, practically no subjects failed to report the additional stimulus.

On average, 25% of subjects reported that they had seen nothing other than the cross in the critical trial. More surprisingly, when the additional stimulus was at fixation, between

⁵ My criticism could easily be adapted to apply to other influential experimental work on inattention blindness – for example Simons and Rensink's (1999) argument that we are blind to unattended dynamic events.

⁶ The point of fixation is the point at which vision is literally most focussed. It's dictated by the relative positions of the eye and external stimuli: the point centred on the fovea, at a certain distance, is the point of fixation. You can attend covertly away from fixation (Posner, 1980).

60% and 80% of subjects reported that they had seen nothing other than the cross. To explain this, Mack and Rock suggest that attention was actively inhibited at fixation. Since it's difficult to direct your visual attention away from fixation, tasks which require you to do this inhibit attention at fixation. To support their interpretation, Mack and Rock used an amended procedure designed to inhibit attention to peripheral locations. On the trials prior to the critical trial, attention was attracted to a peripheral location, by a stimulus irrelevant to the task. This seems to have led to inhibition of attention to the location of the irrelevant stimulus: 50% of subjects now reported having seen nothing other than the cross on the critical trial with the additional stimulus at this peripheral location.

Still, these experiments might suggest that you can be visually conscious of something without attending to it in up to 75% of cases, depending on the circumstances. However, Mack and Rock adduce evidence for the view that some residual attention was responsible for visual consciousness in these cases. They reduced the size of the cross presented in their experiments, so that the additional stimulus was presented well outside the quadrants of the cross. With this amendment in place, subjects reported that they had seen nothing other than the cross on 66% of trials. This suggests that, in the earlier procedures, attention to the cross or to its location could spread to the additional stimulus; once the additional stimulus was far enough from the attended area, subjects reported seeing it far less often.

Mack and Rock claim that, if something were to fall both outside the region to which a subject was primarily attending, and in a region in which attention was actively inhibited, the subject would report not having seen anything there in almost 100% of cases. They argue that when something is completely unattended, the subject wholly lacks visual consciousness of it.

Perhaps one could object that Mack and Rock's data don't support a conclusion as strong as theirs. But I want to focus on a deeper problem with their approach. Where the

subject reports seeing nothing other than the cross, Mack and Rock take this to show that she was not visually conscious of anything other than the cross. They assume that reports to the effect that nothing was seen are definitive evidence that nothing figured in visual consciousness. Why should we assume that? The problem is especially pressing given that we're dealing with a question about attention. Reporting on what we see requires us to attend to it. So the failure to report an object of visual consciousness might reflect a failure to attend to the object, rather than an absence of visual consciousness of the object. Suppose you want to defend *Independence*. Then you should be unmoved by Mack and Rock's argument. They beg the question against your view, by assuming that no conscious experience of an object is in place, where no attentive mechanisms of report on the object are in place.

The problem here runs really quite deep. There is no obvious way to test whether *Dependence* or *Independence* is correct, because in potentially probative cases reports and other fairly explicit decisions form our only convincing evidence concerning the character and presence of conscious experience. We're not limited to asking subjects to describe their experiences. For example, Mack and Rock did not only ask their subjects whether they had seen anything other than the cross. They also asked them to identify what they had seen from among an array of stimuli, or – if they reported having seen nothing – to guess which of these stimuli had been presented. In general, we can ask subjects to make decisions which reveal what they experience consciously, rather than simply asking them to describe their experiences. But the evidence from such decisions is as problematic as the evidence from fully explicit verbal reports. Mack and Rock found that subjects who reported seeing nothing were also unable to identify the stimulus which had been presented. But this failure might reflect either (i) that the subjects were not visually conscious of the stimulus, or (ii) that, though subjects were conscious of the stimulus, they did not attend to it in the way required

for this consciousness to form the basis for a reliable decision. To assume that (i) is the correct interpretation is to beg the question.

On the other hand, there's no obvious way to argue for interpretation (ii) either, because without reports or fairly explicit decisions we lack compelling evidence for the presence of consciousness. We might try looking beyond such decisions and reports for behaviour which reveals conscious awareness. But other sorts of behaviour aren't sufficient to distinguish between the presence of conscious awareness and the presence non-conscious information-processing. For example, Mack and Rock found that the subjects who reported seeing nothing other than the cross were nevertheless affected by the additional stimulus, in ways which showed up in their behaviour. Where the additional stimulus was a word (e.g. 'Flake'), these subjects were significantly more likely than control subjects to use this word to complete a relevant word-stem (e.g. 'Fla-') in a subsequent trial. This is not good evidence for consciousness without attention. The effect may result, as Mack and Rock suggest, from deep but non-conscious processing of unattended stimuli. In general terms, the trouble is that the only behaviour which is distinctive of visual consciousness is behaviour which manifests visual-demonstrative thought. And visual-demonstrative thought involves attention.⁷

⁷ Of course, there is room for sceptical doubt even in the case of behaviour which typically manifests visual-demonstrative thought. In principle, someone might behave in this way without being visually conscious of anything – and so without in fact making visual-demonstrative judgements. But we should set aside this sort of scepticism for practical purposes. It's not because of sceptical doubts that other sorts of behaviour are inadequate to prove that someone is conscious. They're inadequate because we're dealing with difficult cases. For example, we do in general take someone's avoiding obstacles in the street as evidence that she is visually conscious of those obstacles, even though a sceptic might doubt this. The question of what exactly enters her consciousness under conditions of inattention is more difficult, not because scepticism here is

If only we could isolate the neural activity correlated uniquely with visual consciousness, we could establish whether this activity was present in the absence of processes which are correlated uniquely with attention. But of course discovering the neural correlate of visual consciousness would require us first to have established when visual consciousness occurs, and so to have established whether it occurs without attention.

Hilary Putnam, 1981, is so impressed by the difficulties here that he thinks there is, in principle, no hope of discovering an answer to the question whether there is unreportable consciousness, and no hope of discovering which brain activity is correlated uniquely with visual consciousness. Because verbal reports are indispensable as evidence of visual consciousness, Putnam claims that no explanatory grounds could justify a conclusion as to whether the mechanisms of report must be in place in order for someone to be visually conscious. This, he says, is 'a case to be *legislated* rather than fought over' (p.92); there is no fact of the matter, currently eluding discovery by us, about whether you could be conscious yet unable to report it; rather, we can choose which way we want to talk about consciousness.

Some scientists working in this area seem to accept something like Putnam's anti-realist proposal. Victor Lamme, 2005, argues that we should identify visual consciousness with recurrent processing in areas of the brain which process visual stimuli. Recurrent processing is a distinctive kind of neural activity. It can be distinguished from feedforward processing, which moves roughly in one direction from the retina through a hierarchy of visual areas. Recurrent processing, by contrast, involves interactions between areas reached later in the feedforward sweep and areas reached earlier. Recurrent processing can also be distinguished from the processing distinctive of attention: recurrent processing does not entail the kind of neural memory involved in the selection of one stimulus at the expense of

more pressing, but because the question requires us to tease apart different cognitive phenomena which are related in complex ways.

others. Lamme offers no reason for thinking that he has *discovered* an identity between visual consciousness and recurrent visual processing. He cites evidence for the view that recurrent processing, rather than mere feedforward processing, is necessary for visual consciousness. For example, anaesthesia suppresses recurrent processing but not feedforward processing. But he gives no argument for the view that recurrent processing is sufficient for visual consciousness. Rather, he suggests that we 'redefine' visual consciousness as recurrent processing, since recurrent processing is a neurophysiologically identifiable phenomenon which admits of empirical investigation. Thus he legislates against the proposal that visual consciousness depends constitutively on the mechanisms of attention, in addition to recurrent processing.

The best way to resist this sort of anti-realism would be to find a way to argue for, rather than stipulate, a proposal about the relationship between attention and visual consciousness. In Part 3, I'll criticise Ned Block's most recent attempt to argue that visual consciousness is independent of one sort of attention. Then in Part 4 I'll argue that visual consciousness does sometimes depend constitutively on attention.

3. Block on Visual Consciousness Without Cognitive Access

Block, 2007, is acutely aware of the problem I've described, and my discussion so far owes much to his. He's concerned with *cognitive access*, a species of Jamesian attention. Block uses the phrase 'cognitive access' to refer to episodes of thinking about what you see in order to report on it.⁸ These are episodes of Jamesian attention, since reporting on something you see involves concentrating conscious thoughts on it.

⁸ Most often, he uses the phrase to refer to the neural mechanisms of such episodes, but he assumes that conscious episodes are identical with their neural bases (p.482).

Block thinks we can use subjects' reports of their experiences, combined with their discriminatory capacities, to argue that visual consciousness is constitutively independent of this sort of attention, in a way which avoids the problem I've described. If it could be established that visual consciousness is constitutively independent of cognitive access, this would be a step in the direction of an argument for *Independence*. What's more, we'll see that Block's argument presupposes that visual consciousness itself remains static during episodes of cognitive access, rather than undergoing 'focalization' or 'concentration' during them. This suggests, at least, that visual consciousness is constitutively independent of *every* Jamesian focussing of consciousness that occurs during attentive verbal report – a more substantial step in the direction of *Independence*.

Block claims that subjects in some experiments report visual phenomenology which outstrips what they can access cognitively; their reports indicate that they were conscious of details in a stimulus, though they could not say what those details were. He claims that these subjects' performance in discrimination tests provides evidence that their reports are accurate. On this basis, he argues that 'phenomenology and cognitive access are based at least partly in different systems with different properties' (p.494).

By itself, this does not show that visual consciousness or phenomenology is constitutively independent of cognitive access. To complete the argument, Block turns to neurological evidence. Distinct 'coalitions' of neural activation compete for dominance both in visual areas of the brain at the back of the head and in frontal areas associated with access for reports. The coalitions in frontal areas are triggered by dominant coalitions in the back, but losing coalitions in the back may still be strongly activated (Kouider et al., 2007; Deheane et al., 2006). Block argues that this provides a mechanism apt to explain the 'overflow' of visual consciousness beyond what subjects access: while frontal activation is required for cognitive access, strong coalitions in the back of the head are sufficient for visual consciousness, whether or not they are dominant; thus visual consciousness or

phenomenology has a greater capacity than cognitive access. If this proposal about the mechanisms of visual consciousness is correct, then the mechanisms of cognitive access are not 'a necessary part of a neural sufficient condition' on visual consciousness (p.489); neither these mechanisms of access nor the episodes of attention they constitute are among the factors constitutive of visual consciousness; visual consciousness is constitutively independent of one important sort of attention.

As evidence for his claims about 'overflow', Block cites experiments by George Sperling, 1960. I think Block's interpretation of these experiments underestimates the indeterminacy of visual consciousness. Once we take proper account of visual indeterminacy, we can see that Block's argument runs into trouble in much the way I described in Part 2: the only visual phenomenology for which Sperling's experiments provide evidence may, for all the experiments show, be accessed; so Block does not succeed in demonstrating that phenomenology overflows accessibility. And without the premise about 'overflow', his argument cannot get off the ground.

Sperling's first experiment was as follows. Subjects were presented for half a second with a grid of between six and twelve letters, with either three or four letters in each row. You can look at a slightly shorter version of this stimulus on Block's website.⁹ Subjects had to maintain fixation at a point in the centre of the grid. When asked to report which letters the grid contained, each subject could accurately report between 3.8 and 5.2 letters on average. These results were unaffected by the size of the grid, so subjects could accurately report fewer than half of the letters in the largest grids. Yet these subjects 'insist that they have seen more than they can ... report afterwards' (Sperling, p.1); they 'said that they could see all or almost all of the letters' (Block, p.487).

⁹ <http://www.nyu.edu/gsas/dept/philo/faculty/block/demos/Sperling320msec.mov>

In subsequent experiments, trained subjects heard a tone immediately after the visual stimulus ceased. The pitch of the tone indicated which row of the grid they should report – a high tone for the top row, a low tone for the bottom, or a middle tone for the middle row. They were able to report which letters were in the indicated row with a high degree of accuracy – each subject getting between 75% and 100% right after practice, even when faced with the largest grid. This demonstrates that subjects processed information about the specific shape of almost every letter in the grid, even though they could access much less information when their attention was not cued. They must have processed information about the specific shape of almost every letter in the grid, because even once the stimulus was gone they could access that information if appropriately cued.

Block makes a stronger claim. He claims that Sperling's subjects enjoyed conscious, phenomenal experience as of the specific shapes of all or almost all the letters. He claims that this conscious perceptual content was inaccessible, in the sense that subjects were unable to report the specific shapes or identities of nearly half the letters. In that sense, he says, visual consciousness 'overflows' accessibility.¹⁰

¹⁰ As evidence for 'overflow', Block also cites experiments by Landman et al., 2003. These combine a Sperling-style paradigm with a 'change blindness' task: when uncued, subjects could not accurately compare the size and orientation of some of the rectangles in two serially presented grids; when they were cued to the position of a rectangle in each grid, with the cue appearing on a grey screen between the presentations of rectangles, their performance improved. Block takes this to show that subjects were conscious of some rectangles' sizes and orientations, prior to and independently of accessing them. Despite the name, what this task tests for is not awareness of a visible change, but rather a capacity to notice a *difference* – a complex fact about the stimuli, rather than a feature of any individual stimulus (cf. Dretske, 2004). This introduces a complication I don't want to discuss in detail: on the face of it, the data are compatible with uncued subjects having had cognitive access to the sizes and orientations of all the rectangles in

In what sense were subjects visually conscious of the specific shapes of the letters, according to Block? Well first of all they saw them, and drew on their memory of what they saw in their reports. Onset of the auditory cue was controlled by the same switch which turned off the visual stimulus. So the cue was heard sufficiently soon after the visual stimulus ceased for it to be plausible that subjects were still seeing the recently-departed stimulus when they heard the cue and attended to the indicated row of letters. But Block notes that subjects in experiments similar to Sperling's sometimes report deliberately sustaining 'visual imagery' of the stimulus, after strictly speaking they have ceased to see it. He suggests that, if subjects are not still seeing the stimulus when they hear the cue, they draw on 'visual imagery' that's sustained after the stimulus has disappeared. And he counts this as a kind of visual consciousness (p.532).

Block is going on two sets of data: subjects' performance in the task of identifying letters in cued rows, and their verbal reports. The first set of data demonstrates that information about almost all the letters' specific shapes was processed independently of subjects' access to it, but it does not demonstrate that the subjects were *conscious* of these specific shapes independently of access. As we saw in Part 2, evidence that information was processed in your visual system is not necessarily evidence that you were conscious of that information. Block argues that the relevant information was processed in higher-level vision, rather than merely processed in early vision or simply recorded at the retina. But that does not demonstrate that subjects were conscious of this information. Precisely what's at issue here is whether higher-level vision is sufficient for visual consciousness, independently of the mechanisms of cognitive access.

both stimuli, and lacked only the capacity to compare them. In any case, my criticism of Block's discussion of the Sperling experiments applies equally to his discussion of the Landman et al. experiments.

A great deal turns, then, on Sperling's subjects' reports of their experiences. Unfortunately we don't have any quantitative data about these reports. But let's accept Block's claim that subjects reported seeing all or almost all of the letters. One option here would be to deny that we can take these reports at face value. That's what I proposed concerning the reports of Mack and Rock's subjects. But this case differs significantly from that one. Mack and Rock's subjects said that they had seen nothing other than the cross. Presumably, they reflected on their visual experience and found nothing there other than the cross. This admits of two explanations of roughly similar simplicity: (i) they were visually conscious of nothing other than the cross; (ii) though they were visually conscious of the additional stimulus, they failed to notice this. These explanations are of roughly similar simplicity, because (ii) requires us to accept only that noticing what we see is mediated by limited-capacity attentional mechanisms. The case we're dealing with now is different. The idea would be that Sperling's subjects did not really (consciously) see all or almost all of the letters; rather they thought they saw that many, but in fact they saw only a maximum of 5.2 on average. If Sperling's subjects are wrong about their visual experiences they must have made a positive error about them, finding something there which was not in fact there, rather than simply failing to find something which was there. This requires a more elaborate explanation.

Some authors do suggest that Sperling's subjects might think they see more than they actually see (Dehaene et al., 2006). And others suggest a way of explaining how we might come to think of ourselves as visually aware of every part of an array in front of us, even though we're visually aware of only the portions of the array to which we attend (O'Regan and Noë, 2001). The explanation is that, in order to discover whether you're seeing something, you attend to it, thereby becoming aware of it; this can give rise to the illusion that you were aware of the thing all along. The trouble with this line of response to Block is that it presupposes an interpretation of the 'inattention blindness' experiments (Block, 2007,

p.493). Only someone already convinced that visual consciousness requires attention would take such experiments as evidence that we systematically think we see things which we don't in fact see. We can't just make that assumption against Block's argument to the contrary. And without that assumption there's no reason to prefer this more complicated explanation to the simpler explanation that takes Sperling's subjects' reports at face value.

So let's take their reports at face value: Sperling's subjects could see all or almost all of the letters. Now Block takes this to show that their conscious visual experience contained more specific detail than they could access:

[A]lthough one can distinctly see all of the 9-12 objects in an array, the processes that allow one to ... identify the specific shapes are limited by the capacity of 'working memory,' allowing reports of only about 4 of them.

(Block, 2007, p.487)

But why should we accept that the information of which these subjects were conscious outstrips the information they were able to report? As I've said, their performance at the discrimination task demonstrates that they processed information about the specific shape of almost every letter. But why should we take it that they were conscious of this information? An alternative proposal is that subjects were conscious of all the letters, just as they said, but that they were not conscious of the specific shapes of all the letters. Rather, they were conscious of some of the letters in a less determinate way, a way which matches the specificity of the reports they could offer. If they were like me they could report nine characters arranged in a grid, and they could report the identities of only about three or four of those characters. Similarly, the alternative proposal asserts, their visual experience should be characterized as an experience as of nine characters arranged in a grid, only three or four of them appearing sufficiently determinately to appear as a specific letter.

On this alternative proposal, the effect of cuing was not to allow access to a subset of the information of which Sperling's subjects were conscious. Rather, the effect was to alter

their conscious experience such that some information became more determinate in it. This effect might have occurred either in a continuing episode of conscious seeing, or in an episode of sustained conscious visual imagery. Note that this interpretation doesn't require any *retroactive* effect of attention on conscious experience; attention made information more determinate in whatever episode of visual consciousness was ongoing at the time. According to this alternative interpretation, visual consciousness did not outstrip cognitive access. Rather, subjects were conscious of the items in the array with a degree of determinacy which precisely matched the determinacy of their cognitive access. If this interpretation is as plausible as Block's, his argument for 'overflow' fails.

It's important to distinguish the relevant notion of indeterminate representation from some different, neighbouring notions. To represent something indeterminately, as I'm using the phrase, is to represent it as instantiating a determinable property, without commitment as to which determination of that determinable it instantiates. Roughly, property *A* determines property *B* where to have *A* is to have *B* in a specific way.¹¹ For example, the property of

¹¹ This will suffice for present purposes, though it's not precise enough to capture the notion of determination. For example, if there's a conjunctive property of being red and square, then perhaps to have this property is to have the property of being red in a specific way. But intuitively this conjunctive property does not determine being red. The intuitive difference is that anything we add to being red, in order to capture what it is to be crimson, will itself essentially involve being red; being square does not essentially involve being red (cf. Prior, 1949). See Funkhouser, 2006, for a recent attempt to define determination precisely. The notions of determination and indeterminate representation resist formalisation in modal logic. Yablo, 1992, p.252, suggests the following necessary condition on a property *P*'s determining a property *Q*:

- (i) necessarily, for all *x*, if *x* has *P* then *x* has *Q*; and (ii) possibly, for some *x*, *x* has *Q* but lacks *P*.

Intuitively, this condition is probably not sufficient for determination (cf. Yablo, footnote 23). If it were, then for example a conjunctive property would determine each property conjoined.

being crimson and the property of being scarlet each determine the property of being red; being a pentagon determines being a polygon. If you represent something as red without commitment as to which determination of being red it has, or as polygonal without commitment as to which determination of being polygonal it has, you represent that thing indeterminately. Likewise, a representation's being more or less determinate consists in the relative specificity of the most specific determinable property it represents.

So an indeterminate representation need not be a *misrepresentation*. For example, whereas a misrepresentation of a navy blue shirt might represent it as royal blue, an indeterminate representation of the navy blue shirt might tell you that it's blue, without telling you more specifically what colour it is. The mark of indeterminacy in a representation is that the representation is hard to falsify, rather than that the representation is false: the more indeterminate your representation of the shirt's colour is, the wider the range of colours the shirt could have consistent with your getting it right. If you represent the shirt as royal blue, you'll be wrong if it's actually navy blue. If you represent it simply as blue then you'll be right if it's navy blue and wrong if it's red, say.

This sort of indeterminacy is also distinct from Quinean indeterminacy about reference. Quine's thought was, very roughly, that there is no principled reason for saying that a cognitive response represents one aspect of what stimulates it rather than another; the response does not determinately represent any aspect of the stimulus. By contrast, the indeterminacy to which I'm appealing falls within the scope of the representation: a state that's determinately a representation of something as blue represents that thing as having a determinable property; in that sense, the representation is indeterminate between the thing's being royal blue and its being navy blue.

Visual indeterminacy – in my sense – should also be distinguished from what we might call *visual vagueness*. A representation is vague where its extension is not fixed precisely, so that in borderline cases there may be no determinate fact about whether it

represents things as they are. Suppose an object *O* has a colour somewhere on the border between red and pink, and suppose you experience *O* as red. Perhaps, in some such situation, there might be no determinate fact about whether your experience is veridical. By contrast, a representation is indeterminate in my sense where a certain range of possible circumstances falls in its extension. If you experience *O* as red, your experience is veridical where *O* has any colour on the red range of the spectrum. This is compatible with, but does not require, vagueness in the representation. That is, it's compatible with, but does not require, there being no precise fact about where the limits of the relevant red-range lie.

Finally, my claim that visual consciousness is indeterminate should be distinguished from the claim that there is content in visual consciousness which has no semantic value, vague or otherwise. My claim is consistent with (but does not require) the thought that visual content can be exhaustively characterised in semantic terms.

Perhaps further alternatives to Block's interpretation of the Sperling experiments could be framed in terms of these other ideas. But my claim is just that subjects might have experienced some letters in the grid as having determinable shapes, without commitment as to which more determinate shapes they had. For instance, a subject might have experienced an 'F' as having a shape defined by simple lines, without representing it as having specifically the shape of an 'F'.

On some ways of thinking about visual consciousness, this alternative is unavailable. Consider Hume:

'tis confest, that no object can appear to the senses; or in other words, that no impression can become present to the mind, without being determin'd in its degrees both of quantity and quality. The confusion, in which impressions are sometimes involv'd, proceeds only from their faintness and unsteadiness, not from any capacity in the mind to receive any impression, which in its real existence has no particular degree nor proportion.

(Hume, 1740/1978, I.1.vii)

On one reading, Hume claims that visual experience of a property is always experience of a maximally determinate property; you never see something as polygonal, say, without seeing it as the maximally determinate shape in virtue of instantiating which it instantiates the determinable, polygonal. There are certainly other ways to read the passage, but if Hume did claim this, he had respectable reasons for doing so. For Hume, visual awareness is awareness of impressions which resemble the objects you see – that is, impressions which share intrinsic properties with those objects. Impressions are mental items. Like ordinary objects ordinarily conceived, impressions instantiate determinable properties only where they instantiate determinations of them. Your impression is polygonal only if it's polygonal in a determinate way – rectangular, for example. And in being aware of an impression, you're aware of all its properties. So when you see an object in the environment around you, fully determinate properties of the object are apparent to you, since fully determinate properties of the corresponding impression are apparent to you.

On this score at least, we should not think of visual consciousness in Hume's terms. Undeniably, it's tempting to assume that we always see maximally determinate properties. Perhaps this is because of our tendency to reify visual experience, to confuse determinacy in experience with the determinacy of its objects. But whatever the temptation, there are decisive empirical reasons to accept that visual consciousness in fact represents things indeterminately, including with respect to their shapes.¹²

¹² There is also a substantial philosophical tradition which emphasises perceptual indeterminacy. It's striking that this tradition is bound up with a view of the relationship between consciousness and attention which is diametrically opposed to Block's. For example, Leibniz held that ordinary perceptual consciousness consists in apperception of your own internal representational states. By 'apperception', Leibniz means 'attention to what is in us' (1714/1989, §4), and he works with a notion of attention similar to Block's notion of cognitive access: to attend to something is to concentrate higher cognition on it, or on a representation of it, either because of your antecedent

The most general reason lies in the limited and varying *resolution* of vision. For example, the spatial resolution of a representation is given by the maximum spatial frequency it can detect. Spatial frequency consists in the rate of change of a phenomenon across space. In Figure 2, spatial frequency increases from left to right, in that the rate of change from light to dark and back increases from left to right. Where visual spatial resolution is higher, you can see more details of a scene; new edges become visible, and you can see more exactly where those edges are, what orientations they have, and so on. Visual spatial resolution is higher, the closer you are to what you see. It's also highest at the fovea and progressively lower away from the center of the eye. But it's limited even for foveated objects: even when you look right at something, there are spatial frequencies across it which you cannot see (DeValois and DeValois, 1998). In Part 4, we'll see that attention increases visual spatial resolution.

Given these facts about spatial resolution in the visual system, veridical visual consciousness must be indeterminate. The visual system does not detect high spatial

interests or because it grabs attention exogenously (1704/1981, §54). So for Leibniz the machinery of consciousness is identical with the machinery of something like cognitive access.

Leibniz also held that the determinacy of perceptual consciousness varies with apperception. For example, when you attend to the sound of the sea, your overall experience is made up of 'minute perceptions' of the sounds of individual waves; because these minute perceptions are not individually attended, they 'constitute that *je ne sais quoi*, those flavours, those images of sensible qualities, vivid in the aggregate but confused as to the parts' (1704/1981, §54). For Leibniz, confused representations can be contrasted not only with vivid representations, but also with clear representations, and a clear representation is defined as a representation which is sufficient for one to identify what's represented (1704/1981, II.xxix). So each part of the sound of the sea is experienced, but experienced in a way insufficiently determinate for one to identify it, because the overall sound is the object of attention. We might compare the individual letters in Sperling's grid.

frequencies; some changes across space are invisible, making maximally determinate orientations, locations and shapes invisible. If visual consciousness represented such determinate properties, it would practically never do so veridically. In fact, conscious vision represents determinable spatial properties. Where spatial resolution in the visual system is low, visual consciousness of spatial properties is correspondingly indeterminate.¹³

Various further features of visual processing prevent subjects from seeing details of a stimulus, under specific circumstances. One striking example is the *crowding effect*: when a stimulus is presented in the periphery of a subject's visual field, surrounded by other slightly different stimuli, the subject is sometimes unable to identify the specific character of the stimulus. For example, Herman Bouma, 1970, found that subjects can see but not identify letters under these conditions. A natural interpretation is that these subjects see determinable properties, such as the property of having a letter-like shape, rather than determinations of them, such as the property of having an 'F'-shape, say. If you look at Figure 3, you might experience the effect for yourself. This effect occurs within visual experience, not just in our capacity to access experiences of fully determinate properties. The mechanisms of the effect lie in areas of the visual cortex which are demonstrably involved in allowing us consciously to see as we do, perhaps even as early as V1 (Levi, 2008).

¹³ We could accommodate the facts about spatial resolution using an extensionally equivalent alternative to visual indeterminacy. For example, limited-resolution conscious vision might avoid error by qualifying its contents as *approximate*; we might experience things as having properties which approximate certain maximal determinations, rather than experiencing things as having determinable properties. As it happens, I think no such alternative to visual indeterminacy assigns content to visual consciousness in a principled way, either from the point of view of phenomenology or from that of an interpreter. But the arguments of this paper don't turn on a distinction between visual indeterminacy and extensionally equivalent alternatives to it.

I don't know whether Sperling's subjects saw some letters in the grid in resolution too low to identify them. Nor do I know whether they suffered the crowding effect. For one thing, genuine crowding may occur only for peripheral stimuli, and some of the letters which Sperling's subjects couldn't identify may have been at foveal locations. But crowding is one of a family of phenomena involving similar but statistically distinguishable effects – and so perhaps involving different mechanisms; see Levi, 2008, for a review. Some of these effects occur at foveal locations, and some of them are mitigated by attention. Together with the data about spatial resolution, these effects are empirically relevant to the Sperling experiments. They raise the possibility that subjects might have experienced letters indeterminately where they were not the focus of attention.

Block suggests that cognitive access might represent Sperling's letters only 'under a general concept like “alphanumeric character”' (p.487). My objection is that visual consciousness might have a parallel indeterminate character, so that there is no mismatch between the content of visual consciousness and the content of cognitive access. In his published replies to criticism of his paper, Block draws a distinction between 'generic phenomenology' and 'specific phenomenology'. He applies this distinction as follows:

For the Sperling experiment, the relevant generic/specific difference would be that between a phenomenal presentation *that there is* an array of alphanumeric characters and a phenomenal presentation of specific shapes of all or most of the items in the array. I argued that ... there was specific phenomenology involving all or almost all of the items, as well as generic phenomenology.

(Block, 2007, p.531)

Block marshals several arguments for his claim that subjects enjoyed 'specific phenomenology'. I address the key points below.

Block cites subjects' reports that they experience all or almost all of the items in the array. Such reports do not support the conclusion that subjects enjoy 'a phenomenal

presentation of specific shapes of all or most of the items in the array'. They don't support this conclusion, because you might be conscious of an item without being conscious of its specific shape. Perhaps you must be conscious of some features of a thing, in order to count as conscious of that thing. But why suppose that you must be conscious of something's specific shape, in order to be conscious of it? You might be conscious of a letter 'F', say, as having a certain general sort of shape, without being conscious of it as an 'F'. This involves more than Block's 'generic phenomenology' – more, that is, than being aware that there is an array of alphanumeric characters – because it involves awareness of the particular item. But it involves less than Block's 'specific phenomenology' – less, that is, than being aware of the specific shape of the particular item you see. Block attempts to divide conscious visual content into two categories, the specific and the generic. In this way, he artificially excludes the possibility of visual content with a degree of specificity that falls between his two options. And the artificially excluded contents fit subjects' reports of the Sperling experiments just as well as Block's 'specific phenomenology' fits those reports. So Block takes the experiments to favour his view that visual consciousness is independent of cognitive access only because he casts the possibility of indeterminate visual content in terms of a false choice.

This same false choice is central to some of Block's other arguments that Sperling's subjects enjoy 'specific phenomenology'. Block cites evidence that subjects can report accurately whether they are conscious of a whole array, or conscious only of part of it. Subjects report seeing a whole 12-square grid if and only if they are also able accurately to report which of the squares have dots in them (Loftus & Irwin, 1998; Brockmole et al., 2002). This suggests that we should take at face value subjects' reports that they are conscious of all or most of the items in Sperling's array. And it suggests that Sperling's subjects enjoy phenomenology more specific than Block's 'generic phenomenology', given that generic phenomenology doesn't suffice for awareness of any individual item in the array.

But it doesn't suggest that Sperling's subjects enjoy 'specific phenomenology' of all or almost all of the items in the array. It doesn't suggest this, because it's silent on the question of whether subjects are conscious of the specific shapes of the items in the array.

Block notes that 'subjects report no ... phenomenological shift' when they shift their attention to the specific items on which they report. Now my alternative interpretation of the Sperling experiments does require that subjects' phenomenology shift, in the sense that they become conscious of more specific details of the items to which they shift their attention. But it's not at all obvious that this shift is something subjects would, or could, report. Suppose you experience some of the items in Sperling's array with a shifting degree of determinacy. Suppose that you're first conscious of a letter-like shape in the bottom corner of the grid, and that when you shift your attention to that shape, you become conscious of the shape as an 'F'. How could you distinguish this from a case in which you were conscious of the shape as an 'F' all along, though you weren't attending to this aspect of it? Attention to what you see is required if you're to report on changes in the determinacy of your visual consciousness. But if attention to what you see *effects* changes in the determinacy of your visual consciousness, keeping track of the changes will be difficult at best.

Block takes it that the shift in phenomenology should be introspectively obvious, because he takes the shift to be between 'generic phenomenology' and 'specific phenomenology'. That sort of shift might well be obvious, given that it would be a shift between, first, content which might be captured in a proposition about the array which makes reference to no specific item, and subsequently, content which might be captured in a detailed picture of all the items in the array. But the choice between 'generic phenomenology' and 'specific phenomenology' is false, and the shift need not be so dramatic.

Block also appeals to a different sort of report about experiences of the Sperling experiment. Subjects claim, he says, to be 'reading their answers off of the visual impression that was in existence before the cue'. Here again, we have no formal data about exactly what

subjects say. But let's accept that they say something along these lines. Note that subjects do *not* say that they had a detailed visual impression which was unaltered by their shifting attention. To be sure, they were visually conscious of the whole grid before their attention was cued. And to the extent to which talk of visual impressions is appropriate here, subjects read their answers off their visual impressions. But this does not entail that their visual impressions were unaltered by the acts of attention involved in reading off their answers. Perhaps subjects' impressions became determinate with respect to specific shapes only where they cognitively accessed those specific shapes. If I ask you to form a visual image of your kitchen, and then ask you to report the specific shape of your stove, the image you enjoy will change as you shift your attention to the stove. The image will change in that, *inter alia*, you will represent the stove more determinately. Still, you were 'reading your answer off of the visual impression'.

The idea that Sperling's subjects read their answers off an image does not support Block's claim that there was 'specific phenomenology' prior to cognitive access to specific shapes. Nor does anything else about the Sperling experiment support the claim that visual consciousness 'overflows' cognitive access, rather than precisely matching the content of cognitive access. I haven't ruled out the possibility that subjects were conscious of details they were unable to report. But there is equally no reason to think that they were so conscious. Where reports are non-specific, it may be that consciousness matches them. So neither set of experiments favours Block's claim that consciousness overflows cognitive access, any more than its denial. It's only because Block underestimates visual indeterminacy that he seems to avoid the problems described in Part 2.¹⁴

¹⁴ Van Gulick (in Block, 2007) suggests, as I have been, that visual consciousness might be indeterminate and so match cognitive access in its content. But Van Gulick doesn't spell out indeterminacy as I have (Grush's criticism is probably closest to mine), and Block's response is

4. An Argument for *Dependence*

The foregoing leaves us with the *impasse* I identified in Part 2. Reports and other visual-demonstrative judgements form the only good evidence concerning visual consciousness. For this evidence to be present, attention must be present. So we cannot argue for *Dependence* by appeal to the absence of this evidence, on pain of begging the question. Conversely, Block cannot establish that things figure in visual consciousness independently of whether you access them, since reports demonstrate that a feature of the stimulus figured in visual consciousness only where, or to the extent to which, the subject attended to that feature.

I think this *impasse* results from the debate's being framed in terms of an overly demanding criterion for constitutive dependence. Mack and Rock try to show that attention

independent of his specific/generic distinction: Block says there's 'experimental evidence to the contrary'; 'cognitive access seems to be more of a binary phenomenon than Van Gulick supposes'. The evidence Block cites doesn't tell against the possibility I've raised – that cognitive access matched visual consciousness, in that both visual consciousness and cognitive access represented some items in Sperling's array without representing them as specific alphanumeric characters. This evidence concerns the 'attentional blink' (Sergent & Dehaene, 2004). Subjects are asked to identify two sets of letters (e.g. XOOX and FIVE) from among a series of brief four-letter presentations. Given certain delays between the target presentations (especially 0.3 seconds), subjects fail to report the second target. When they're asked to rate, on a scale from 1 to 10, how visible the second target was, they rarely choose intermediate values. This suggests that whether or not the attentional blink prevents access to a particular letter is an all-or-nothing matter. But *that* might be an all-or-nothing matter, even if the shapes cognitive access represents vary in determinacy.

is necessary for visual consciousness – which entails that attention is part of a minimal sufficient condition on *every* episode of visual consciousness. Block tries to show that cognitive access to details of a scene is not necessary for visual consciousness of them; then, by appeal to his hypothesis about separate neural mechanisms, he argues that cognitive access is *never* part of a minimal sufficient condition on visual consciousness. These theses are certainly interesting. Perhaps they hold the promise of telling us whether or not being visually conscious essentially involves attending or accessing. But we can settle the constitutive question about *Dependence* and *Independence* without arguing for anything so bold. Minimal sufficient conditions on a token episode of conscious vision need not be conditions on *every* episode of conscious vision. I'll argue for *Dependence* by arguing that, in some token cases, your attending is part of a minimal sufficient condition on your seeing as you do.

In light of Parts 2 and 3 above, I'm agnostic about whether frontal neural processes or cognitive access ever fall in the set of factors constitutive of visual consciousness. But episodes in the visual cortex, for example recurrent processes involving V4, are a good candidate for being among this set of factors. Changes in these processes are implicated in changes in the course of conscious visual experience (Lamme, 2005), which suggests that they form part of a minimal sufficient condition on that experience (Block, 2005).

It's controversial how *broad* the constitutive basis of visual consciousness might be. Hurley and Noë, 2003, argue that some changes in the phenomenal character and content of conscious vision require only change in parts of a dynamic neuro-environmental system that lie outside the skull. Others argue that the phenomenal character of conscious vision should be identified with the qualitative character of the objects we see (Campbell, 2002; Martin, 2002). And Williamson, 2006, argues that, to do justice to the causal powers of cognitive states, we must think of those states as constitutively involving the states of affairs they represent. Any or all of these ideas might lead us to include in the set of events constitutive

of visual consciousness the whole of a causal system reaching from the visual cortex to events in the environment around a perceiver. Attention might well figure in this causal system, and so also figure in this broad constitutive basis. However, the constitutive connection between attention and visual consciousness need not be hostage to these controversial issues. So, for the sake of the strongest possible argument for *Dependence*, I'll allow that the events constitutive of visual consciousness may occur exclusively within the visual cortex.

An argument for *Dependence* will need the following form. First I need to identify an episode of attention, an episode of selective 'focalization, concentration of consciousness'. Then I need to show that this episode falls within the minimal set of factors sufficient for the character and content of an episode of visual consciousness. The evidence I'll appeal to concerns the spatial determinacy of conscious vision, and comes from work by Yaffa Yeshurun and Marisa Carrasco (1998).

As I've said, visual spatial resolution is highest at the fovea, and progressively lower away from the centre of the eye. Early in post-retinal visual processing, information about the scene you see passes through varying spatial filters, attuned to various spatial frequencies. As the locations of these filters get further from the fovea, they're attuned to progressively lower spatial frequencies: higher-frequency detail goes unprocessed (De Valois and De Valois, 1988).

A texture segregation task is a task in which you have to report whether a stimulus contains a section in which lines have a unique orientation, relative to the lines in other sections of the stimulus (Figure 4). In general, seeing with higher spatial resolution is useful for texture segregation. So in general texture segregation becomes easier as the unique texture approaches foveation. Yeshurun and Carrasco used texture-segregation tasks to investigate how attention interacts with visual spatial resolution (Figure 5). Holding fixation fixed, they cued conscious attention by flashing a small bar immediately above a location in

a texture-segregation stimulus. When they cued attention to a peripheral location containing a texture-segregation target, they found that attending to that location made the task easier.

Now this result is consistent with *Independence*. In attending to the cued location, you might simply focus your thoughts or cognitive access on it, and in that way become more aware of its distinctive features. Attending might merely bias your thoughts and decisions to a location, without affecting visual consciousness itself. This would be an example of Jamesian attention, an example of selective 'focalization, concentration of consciousness', but the focalization would take place at the level of thought, not at the level of visual consciousness.

However, there are some texture segregation tasks in which seeing with too high a spatial resolution makes it harder to identify the unique texture. As Yeshurun and Carrasco put it, '*when a more global inspection of the display is required, for example when one is appreciating an impressionist painting, moving closer is not the optimal strategy*' (p.74). As a result, there are some texture segregation tasks at which we're worse as the unique texture approaches foveation. Yeshurun and Carrasco exploit this phenomenon to argue against the view that, in the cases they studied, attention merely affected subjects' decision-making procedures, leaving the visual signal static. They argue that attending enhances spatial resolution in the visual signal itself, by mimicking the effect of foveation, by effectively making the spatial filters for a location more finely tuned. In my terms, they tested the hypothesis that visual consciousness was *not* static through changing attention, that attention increased the spatial resolution of visual consciousness itself.

In a task in which texture segregation is harder where spatial resolution is too high, they found that cuing attention to the unique texture *impaired* performance, when that texture was at or near foveation. This result cannot be explained in terms of the idea that attention simply improves your access to what's represented in a static visual consciousness. According to that proposal, attention should only improve the decisions you make, by

devoting more decision-making resources to the relevant location. Attending to the unique texture should never diminish performance at texture segregation. To explain the data, we must accept that attention increases the spatial resolution of visual consciousness, replicating the effect of foveation.

Yeshurun and Carrasco characterise the relationship between attention and visual consciousness in causal terms. But their evidence also reveals a closer, constitutive connection between conscious, Jamesian attention and visual consciousness. In the cases they investigate, Jamesian attention is part of a minimal sufficient condition on the phenomenal character and content of an episode of visual consciousness. As a result of attentional cuing, visual content becomes more determinate. This increase in determinacy is a selective 'focalization, concentration of consciousness'. From the point of view of an interest in conscious attention, Yeshurun and Carrasco's discovery is that Jamesian attention, the selective 'focalization, concentration of consciousness', sometimes takes the form of an increase in the determinacy of conscious vision. As a result, the episode of attention is among the set of factors sufficient for the changing content of visual consciousness, and such that no subset of this set is so sufficient. Someone attending in this way is, necessarily, also enjoying a visual experience with a certain changing content. And nothing short of attending in this way would suffice, under the circumstances, for enjoying this experience.

This constitutive connection is reflected in the fact that the neural basis of visual consciousness here includes the neural basis of attention. Consider one candidate neural basis. It's well-established that attention produces heightened selective responses in V4 (Moran and Desimone, 1985; Luck et al., 1997). These responses are a good candidate for being at least part of the neural basis minimally sufficient for the 'focalization, concentration of consciousness' involved in seeing with a greater spatial determinacy (cf. Yeshurun and Carrasco, p.74). For that same reason, they're a good candidate for being at least part of the neural basis minimally sufficient for the changing content of visual consciousness. Whatever

is minimally sufficient for the episode of attention is also part of what's minimally sufficient for the changing episode of visual consciousness.¹⁵

I've argued that attention was constitutive of visual consciousness in the cases studied by Yeshurun and Carrasco, 1998. The attention manipulated there was exogenously controlled. Yeshurun et al., 2008, studied the interaction between visual spatial resolution and endogenously controlled attention, directed voluntarily by the subject. Here the results are more ambiguous: attention improves performance at texture-segregation as before, but it never impairs performance. So there's no positive evidence here that endogenous control of attention has the same conscious effect as exogenous control, increasing visual determinacy. The idea that endogenous control has this effect is not *inconsistent* with the data: if endogenous control does increase the spatial resolution of visual consciousness, the absence of impaired performance might be explained by further effects of endogenous control, over and above this increase in determinacy – e.g. sustained processing over a longer period of time.

The argument for *Dependence* does not turn on the question whether endogenous control of attention increases the determinacy of visual consciousness. *Dependence* requires only a constitutive connection between one sort of attention and visual consciousness, so the Yeshurun and Carrasco, 1998, result gives us the resources we need. Furthermore, whether

¹⁵ This proposal is consistent with some recent claims about a neural basis of attention outside the visual cortex. For example, I don't dispute the claim of Armstrong and Moore, 2007, that the *Frontal Eye Field* drives endogenously controlled attention. Rather, following James I treat 'attend' as a success verb: to attend is to achieve a 'focalization, concentration of consciousness'. Accordingly, the neural basis of attention lies in the processes which realise this conscious change. Armstrong and Moore could agree, I take it, that these lie at least partly in the visual cortex. They report that stimulating a monkey's Frontal Eye Field caused increased differential responses to differently oriented stimuli, in neurons in V4.

or not endogenous control *brings about* the conscious effects brought about by exogenous control, voluntary conscious attention is characterised by these effects where exogenous control brings them about. For example, Yeshurun and Carrasco's subjects exercised voluntary attention on the stimuli they saw. They deliberately attended to the stimuli, concentrating consciousness on them in order to identify unique textures. Where attention was also cued exogenously, this concentration of consciousness included an increase in the determinacy of the spatial properties subjects saw. That's true whether or not the deliberate decision to attend was responsible for bringing about this increase in determinacy.

It's possible to replicate the attentional effects which Yeshurun and Carrasco studied, in the absence of attention. If you bring the stimulus closer to your eyes, higher spatial frequencies may become apparent, just as they become apparent when you attend. Perhaps this tells us that attention is not necessary for visual consciousness to have a certain character and content. But it does not count against the hypothesis that visual consciousness sometimes depends constitutively on attention. This is the hypothesis that, in some cases in which attention is responsible for the character of visual consciousness, conscious attention is among the set of factors minimally sufficient for the episode of visual consciousness. The evidence concerning attention and spatial resolution strongly favours this hypothesis.

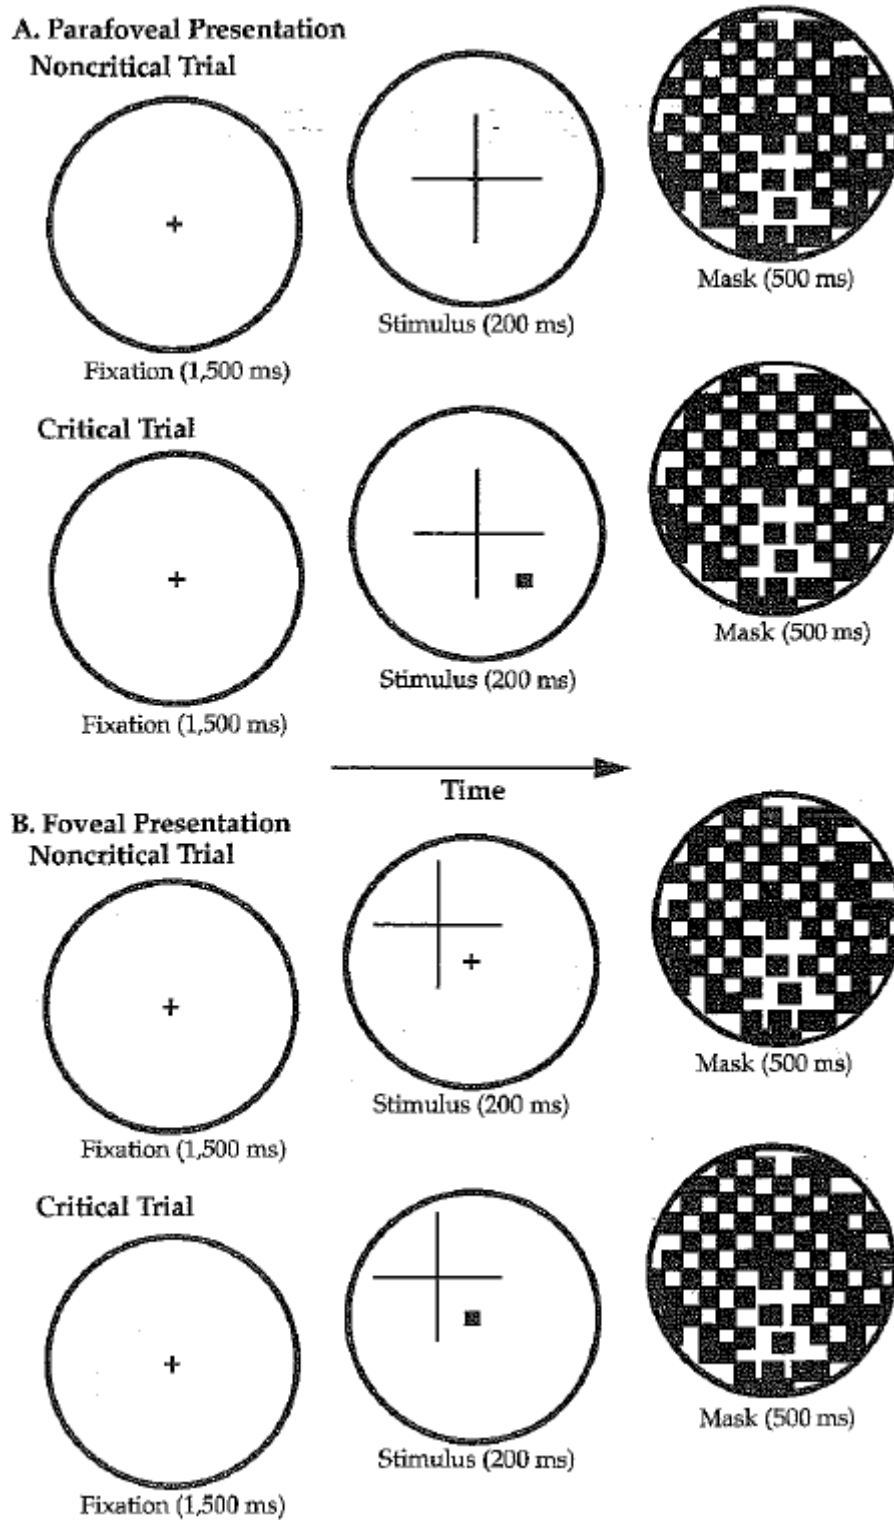
You might insist on using the term 'attention' differently, so that the 'focalization, concentration of consciousness' involved in seeing with increased spatial determinacy does not count as attention. For example, you might insist that attention be defined anatomically, and so that visual-cortical events are not constitutive of it. Or you might insist that attention be defined in terms of higher-order consciousness. As I've said, the term 'attention' is used in various ways. However, the possibility of talking in these different ways does not compromise my claim that the phenomenon defined by James, by appeal to ordinary experience and folk psychology, is constitutive of visual consciousness.

*Department of Philosophy,
University of California, Berkeley*

Total words: 13,279

(Main text and footnotes: 12,224)

Figure 1



Reproduced from Mack and Rock, p.57, by permission of Arien Mack and Oxford University Press, Inc. (copyright 1998).

Figure 2





Figure 3 -- Fixate the dot in box (a), and try to identify the single letter below. Then do the same for central letter of the group, in boxes (b), (c) and (d).

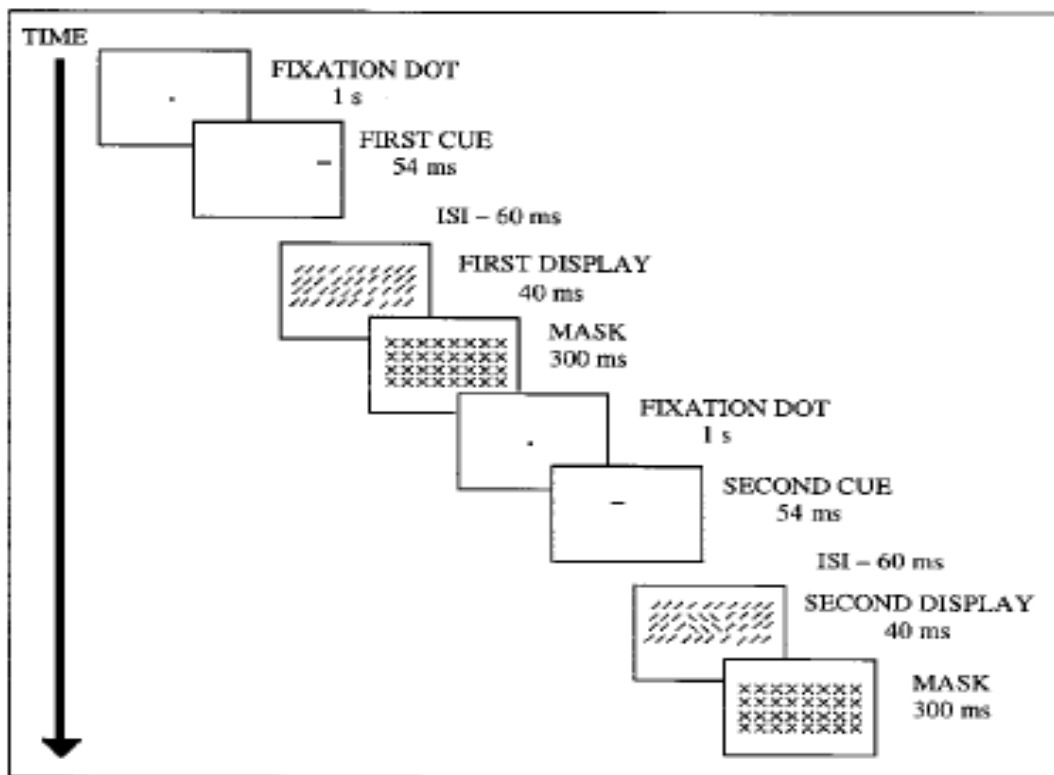


Reproduced from Levi, p.636, by permission of the author and Elsevier (copyright 2008).

Figure 4



Figure 5



Reproduced from Yeshurun and Carrasco, p.72, by permission of Marisa Carrasco and Macmillan Publishers Ltd (copyright 1998).

References

- Armstrong, K. and Moore, T. 2007: Rapid enhancement of visual cortical response discriminability by microstimulation of the frontal eye field. *Proceedings of the National Academy of Sciences*, 104(22): 9499–9504.
- Block, N. 2005: Two neural correlates of consciousness. *Trends in Cognitive Sciences* 9(2): 46-52.
- Block, N. 2007: Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30: 481–548.
- Bouma, H. 1970: Interaction effects in parafoveal letter recognition. *Nature* 226:177-8.
- Brockmole, J., Wang, R. and Irwin, D. 2002: Temporal integration between visual images and visual percepts. *Journal of Experimental Psychology: Human Perception & Performance* 28(2): 315–34.
- Campbell, J. 2002: *Reference and Consciousness*. Oxford: Clarendon Press.
- Davidson, D. 1969: The individuation of events. In N. Rescher (ed.), *Essays in Honor of Carl G. Hempel*. Dordrecht: D. Reidel Publishing Company.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J. & Sergent, C. 2006: Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, 10(5): 204–11.
- DeValois, R. & DeValois, K. 1988: *Spatial Vision*. Oxford: OUP.
- Dretske, F. 2004: Change blindness. *Philosophical Studies*, 120: 1-18.
- Funkhouser, E. 2006: The determinable-determinate relation. *Nous* 40(3): 548-569.
- Hume, D. 1740/1978: *A Treatise of Human Nature*, ed. R. Selby-Bigge / P. Nidditch. Oxford: OUP.
- Hurley, S. and Noë, A. 2003: Neural plasticity and consciousness. *Biology and Philosophy*, 18: 131-168.
- James, W. 1890/1950: *The Principles of Psychology*, ed. G. Miller. New York: Dover Publications.
- Koch, C. 2004: *The Quest for Consciousness: A Neurobiological Approach*. Englewood, Colorado: Roberts and Co.
- Koch, C. & Tsuchiya, N. 2006: Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, 11(1): 16-22.

- Kouider, S., Dehaene, S., Jobert, A. & Le Bihan, D. 2007: Cerebral bases of subliminal and supraliminal priming during reading. *Cerebral Cortex*, 17(9): 2019–29.
- Lamme, V. 2005: Independent neural definitions of visual awareness and attention. In A. Raftopoulos (ed.), *Cognitive Penetrability of Perception*. Hauppauge, New York: Nova Science Publishers.
- Landman, R., Spekreijse, H. and Lamme, V. 2003: Large capacity storage of integrated objects before change blindness. *Vision Research*, 43(2): 149–64.
- Leibniz, G. 1704/1981: *New Essays on Human Understanding*, ed.s P. Remnant and J. Bennett. Cambridge: Cambridge University Press.
- Leibniz, G. 1714/1989: Principles of nature and grace, in his *Philosophical Essays*, ed.s R. Ariew and D. Garber. Indianapolis: Hackett Publishing Company.
- Levi, D. 2008: Crowding—An essential bottleneck for object-recognition: A mini-review. *Vision Research* 48: 635-654.
- Loftus, G. and Irwin, D. 1998: On the relations among different measures of visible and informational persistence. *Cognitive Psychology*, 35: 135–99.
- Luck, S., Chelazzi, L., Hillyard, S. and Desimone, R. 1997: Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology* , 77: 24–42.
- Mack, A. & Rock, I. 1998: Inattention blindness: perception without attention. In R. Wright (ed.), *Visual Attention*. Oxford: OUP.
- Marc-Wogau, K. 1962: On historical explanation. *Theoria*, 28: 213-233.
- Martin, M. 1998: Setting things before the mind. In A. O'Hear (ed.), *Contemporary Issues in Philosophy of Mind*. Cambridge: Cambridge University Press.
- Martin, M. 2002: The transparency of experience. *Mind and Language*, 17: 376-425.
- Martin, M. 2006: On being alienated. In T. Gendler and J. Hawthorne (ed.s), *Perceptual Experience*. Oxford: OUP.
- Merleau-Ponty, M. 1945/1962: *Phenomenology of Perception*, trans. C. Smith. London: Routledge and Kegan Paul.

- Moran, J. and Desimone, R. 1985: Selective attention gates visual processing in the extrastriate cortex. *Science*, 229: 782–784.
- O'Regan, J.K. and Noë, A. 2001: A sensorimotor approach to vision and visual consciousness. *Behavioral and Brain Sciences*, 24: 883–975.
- Posner, M. 1980: Orienting of attention, *Quarterly Journal of Experimental Psychology*, 32: 3–25.
- Prinzmetal, W. and Landau, A. 2008: Dissecting spatial visual attention. In V. Coltheart (ed.), *Tutorials in Visual Cognition*. Hove: Psychology Press.
- Prior, A. 1949: Determinables, Determinates and Determinants (Part I). *Mind* 58, No.229, 1-20.
- Putnam, H. 1981: *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Sergent, C. and Dehaene, S. 2004: Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science*, 15: 720–28.
- Simons, D. and Chabris, C. 1999: Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28: 1059-1074.
- Sperling, G. 1960: The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74 (11, Whole No. 498):1–29.
- Williamson, T. 2006: Can cognition be factorised into internal and external components? In R. Stainton (ed.), *Contemporary Debates in Cognitive Science*. Oxford: Blackwell.
- Yablo, S. 1992: Mental causation. *Philosophical Review*, 101(2): 245-280.
- Yeshurun, Y. and Carrasco, M. 1998: Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*, 396: 72-75.
- Yeshurun, Y., Montagna, B. and Carrasco, M. 2008: On the flexibility of sustained attention and its effects on a texture segmentation task. *Vision Research*, 48: 80–95.