

Color, Qualia, and Attention: A Non-standard Interpretation

Austen Clark, University of Connecticut.

Forthcoming in Jonathan Cohen & Mohan Matthen (eds), *Color Ontology and Color Science*. MIT Press. Revised May 2008.

Abstract. A standard view in philosophy of mind is that qualia and phenomenal character require consciousness. This paper argues that various experimental and clinical phenomena can be better explained if we reject this assumption. States found in early visual processing can possess qualitative character even though they are not in any sense conscious mental states. This non-standard interpretation bears the burden of explaining what must be added to states that have qualitative character in order to yield states of sensory awareness or sensory experience. I argue that the study of selective attention reveals resources that can be useful in that project. Two traditional objects are briefly considered.

Acknowledgments. I thank Jonathan Cohen and Mohan Matthen for all their work organizing the color science conferences and the resulting volume of papers, and for providing very useful comments to me on this particular paper as well. Thanks also go to Valteri Arstila, Murat Aydede, Steven Biggs, and Alex Byrne, who provided comments on versions of this paper, as well as to audiences at the University of Toronto, Harvard University, the University of British Columbia at Vancouver, Simon Fraser University, and Rutgers University, where versions of these ideas were presented.

A state of "phenomenal consciousness" is, minimally, a state of consciousness that has some kind of phenomenal character. Both clauses are open to multiple interpretations, the discussion of which can grow quite heated. For the purposes of this paper, I propose to avoid as many of those debates as possible by confining the discussion to a simple kind of mental state which all sides agree belongs to the genus "phenomenal consciousness", no matter what other species the latter might include. This simple kind might be called "perceptual awareness" or "perceptual experience": episodes in which a subject both perceives something and is aware of what is perceived or of some aspect of what is perceived. Its perceptual origin gives the episode its distinctive phenomenal character, and often that phenomenal character is the feature of which one is aware. Examples include seeing and thereby becoming aware of the red glow of the sky at sunset, or feeling the cool breeze that springs up as the sun goes down. One might also call these "sensory experiences": one experiences--becomes conscious of--some sensible feature of something that one senses.

In these states one both senses something and is aware of some aspect of what one senses. The latter condition is what qualifies the members of this species as "states of consciousness". In virtue of having them one is aware of something: here, specifically, one is aware of some aspect of what one senses. One may, or may not, also be conscious of being in that state. It is possible that one is so absorbed by the sunset--the red glow of the sky, the cool breeze--that one is not also simultaneously aware of *seeing* the red glow, or of *feeling* the cool breeze. We talk of "losing oneself" in various experiences, aesthetic and otherwise, and in such cases one seems to be aware simply of what one

sees, or what one feels; not also of the seeing of it, or of one's own internal state of sensation. So a state of perceptual awareness (as used here) is not necessarily a state *of which* one is conscious. Instead it is a state in virtue of which one is conscious of what one perceives.

Now phenomenal consciousness is a puzzle, and intent study of neuropsychology makes it, if anything, more puzzling rather than less. Among the most startling of incongruities are the phenomena found in the "disconnection" syndromes of blindsight and hemi-neglect. Patients with various kinds of brain damage can pass many of the normal tests for seeing things and having states with phenomenal character, even though they sincerely deny seeing those things, or of being aware of any features of those things. The blindsight subject DB can pass many of the normal tests for seeing shape, orientation, and location, and for discriminating different shapes, orientations, and locations; though if the stimuli in question are presented within his scotoma, DB sincerely denies seeing the stimulus, and denies being aware of any feature of anything located within the scotoma. In pure "unawareness" mode, DB is not aware of any sensible feature of a stimulus presented within the blind field, yet he successfully "guesses" its location, shape, and orientation.

These results seem to drive a wedge between the "phenomenal" and the "consciousness" parts of "phenomenal consciousness." They challenge the widely held assumption that states with phenomenal character are necessarily states of consciousness. More broadly, they pose a puzzle for anyone interested in the relations between perceptual appearance and the awareness thereof. In the first of the two conferences that prompted this volume, Larry Hardin revealed his longstanding interest in this problem. He posed (as an agenda item for conference two) the question "What does the awareness of color add if one can already make chromatic discriminations?"¹ It is a wonderful question, and I propose to address it by examining some of the "disconnection" phenomena. While this essay will certainly not answer the question, I hope the energetic beating of bushes will give us some sense of the surrounding undergrowth, and perhaps flush out some of its interesting and reclusive inhabitants.

I. Blindsight

Blindsight is well known, if not well understood, so I can provide a brief summary and point the reader to fuller descriptions found elsewhere (see Weiskrantz 1986, 1997). Certain kinds of brain damage produce a scotoma, or blind spot, within some portion of visual perimetry. That portion is delimited retinotopically, by azimuth and elevation relative to the point on which the eyes are focused. If the eyes are focused on a fixed point, and stimuli are presented within the scotoma, the patient will, for many classes of stimuli,

1 My notes indicate that the first question Larry Hardin posed was "Why should there be any color experience at all, rather than just information processing?" The question listed above came second, as a followup and elucidation of this first one.

deny seeing any stimulus, and indeed deny being aware of anything at all in that region. Nevertheless, if the patient is prompted to hazard a guess about various attributes of the stimulus, or somehow to indicate a value, sometimes the resulting indicators correlate very highly with the actual values of the stimulus presented. Attributes can include location (indicated by pointing), shape (X vs. O), motion (vertical v. horizontal), and, most surprising of all, color.

The latter deserves some elaboration (see Stoerig & Cowey 1992). Color discrimination is tested through a long series of forced-choice discrimination tasks. Any given series of trials involves just two spectral colors, matched for luminance, which are first named for the subject ('orange' and 'yellow', for example) and presented for examination. Then in a given trial the subject fixates on a central point and one of the two colors, picked randomly, is presented in that subject's scotoma. Since the subject cannot tell whether the stimulus is present or absent, a click indicates when it is time to respond; the task is to "guess" the name of the color that had just been presented. But with a long enough series of such guesses, one can assess the ability of the subject to discriminate different wavelengths. Other tests allow one to compare the spectral sensitivity curves (specifically, threshold sensitivity at different wavelengths) with that of normals (see Stoerig & Cowey 1991). The absolutely stunning result is that some blindsight subjects can indeed discriminate some wavelength differences independently of luminance, at statistically significant levels, even for relatively small wavelength differences (20 nm, the difference between the yellow and the orange stimuli). As Weiskrantz puts it:

The latter capacity--discrimination of colors--presses credulity to the limit, because in those tests--which by their nature were very time-consuming and lasted for several days--the subjects uniformly and consistently denied seeing color at all, and yet performed reliably above chance, even between wavelengths falling relatively close together. Moreover, the fine-grained features of the spectral sensitivity curves of these subjects (carried out, again, by forced-choice guessing) suggested that wavelength opponency, that is, color contrast, was intact (Stoerig and Cowey, 1989, 1991, 1992). The subjects seemed to be able to respond to the stimuli that would normally generate the philosophers' favorite species of 'qualia', namely colors, but in the absence of the very qualia themselves! (Weiskrantz 1997, 23).

In this last sentence Weiskrantz puts his finger on the problem. Behaviorally these subjects pass some of the tests that would normally be taken to indicate that they see the differences between different colors. They pass some of the tests that would indicate that they are sensitive to differences among chromatic appearances. Yet they resolutely and sincerely deny being aware of any aspect of the stimulus presented within the scotoma. So do they have chromatic qualia, or not?

II. The standard interpretation

If you think the answer to this question is "obviously not!", then there is a non-negligible probability that you share an interpretation of the experimental and theoretical literature that I will call the "standard" interpretation. I admit the

label is tendentious; it is "standard" only in the sense that it seems to be a common view in both philosophy and psychology. On this view, qualia and phenomenal properties require consciousness. More precisely, if at time t a subject is not aware of any color at all, then it follows necessarily that at time t that subject does not have any chromatic qualia. For any property to be a quale Q , some subject must be aware of, or conscious of, Q . So showing that DB is not aware of any colors of stimuli in his scotoma settles definitively the question of whether the perceptual states responsible for his discriminations of stimuli therein sport any qualia. They do not and cannot. Similarly, characterizing a property as a "phenomenal" property, or as a property of appearance, implies (according to this view) that some subject is aware of that property. Barring awareness is a way of barring the applicability of all the phenomenal locutions--all the being-appeared-to's, the looking-like F 's and seeming-to-be G 's.

It is easy to produce examples of philosophers promulgating this view, though no such list can suffice to show it is the "standard" view, or even a widely held view. Nevertheless, a few such examples will have to suffice. One classic source is Thomas Nagel. What he calls "phenomenological qualities" are subjective:

A feature of experience is subjective if it can in principle be fully understood only from one type of point of view: that of a being like the one having the experience, or at least like it in the relevant modality. The phenomenological qualities of our own experience are subjective in this way. (Nagel 1979b, 188).

This subjective or "perspectival" character is, according to Nagel, what we are talking about when we say the state is a *conscious* state (Nagel 1979a, 166). John Searle (1997) is a recent advocate of both these claims. Consciousness is a subjective phenomena; it has a "first person ontology" (Searle 1997, 98, 114, 120). Qualia, or characteristics of appearance, cannot be divorced from consciousness:

I myself am hesitant to use the word "qualia" and its singular "quale" because they give the impression that there are two separate phenomena, consciousness and qualia. But of course, all conscious phenomena are qualitative, subjective experiences, and hence are qualia. There are not two types of phenomena, consciousness and qualia. There is just consciousness, which is a series of qualitative states. (Searle 1997, 8-9).

Finally, David Chalmers provides some classic statements of the view that "the problem of explaining these phenomenal properties is just the problem of explaining consciousness" (Chalmers 1996, 4). He says that the words "qualia" or "phenomenal" pick out approximately the same class of phenomena as "consciousness", and that "'To be conscious' in this sense is roughly synonymous with 'to have qualia'..."(Chalmers 1996, 6). Color sensations invariably present qualia, and so color sensations are the "paradigm examples of conscious experience" (Chalmers 1996, 6). There is indeed a robust sense in which to characterize something phenomenally is to characterize how that something is experienced. What I am calling the "standard" interpretation leaves no space at all between the concepts of being-appeared-to and being

aware-of; the former locutions apply only where the latter are ensconced.

Among psychologists the influence of the standard interpretation is often seen in a reluctance to apply certain locutions to the states involved in early, preattentive, perceptual processing. Although DB can (with good though not perfect reliability) discriminate the orientations and shapes of stimuli in his scotoma, it is widely thought that all such discriminations can be explained by the brute mechanical registration and processing of information in early visual channels. Provided DB operates in pure "unawareness" mode, there is no call to say that DB "senses" any such features; that he is presented any sort of appearance by stimuli in the scotoma; that those stimuli have any sort of phenomenal property; or that any of the states that subserve his successful discriminations have any sort of qualitative character. All of these locutions, it is thought, imply that DB is aware of the features, appearances, or characters involved. DB manifestly is not. Ergo any account of the success of those discriminations must avoid any mention of sensation, appearance, or phenomenology. The words are banned. We have instead mere information processing, of the sort that occurs in an utterly unconscious automaton--a telephone switch, a thermostat, or a lap-top computer. Note that this is the contrapositive of the philosophers' version of the standard view. It reads: if one is not conscious of property *P*, then *P* is not a quale, not a property of appearance, not a phenomenal property. We see this view at work in the Weiskrantz quote above: since DB is not aware of the features he successfully discriminates, those discriminations proceed in the absence of the "very qualia themselves".

Now of course everyone has Humpty-Dumpty's right to use a word however they please, but I will argue that insisting on this usage erects a permanent road-block in the way of our understanding phenomenal consciousness. If the argument works, then it will be worth looking at some alternatives that might allow us to skirt around that road-block.

III. Phenomenal Properties

Terms such as "qualia" and "phenomenal" have become loaded and overloaded with multiple and conflicting theoretical connotations, so it is a good idea to start with the simplest and most traditional sense of the word, and confine ourselves to that. "Phainomenon" is Greek for "appearance", and a "phenomenal property", on one traditional reading, is a manner of appearance, a characteristic of how things appear. The "things" in question are, for the purposes of this paper, resolutely confined to the things one perceives; their "appearances" are characteristics of how they (for example) look, or feel, or taste, or (generally) seem. If the mountains look blue in the distance, then "blue" is used in that context to attribute a phenomenal property, characterizing how those mountains appear. We have other "verbs of appearance" that are used, like "looks", to characterize such appearances while remaining neutral on the question of whether those appearances are sustained by reality (see Chisholm 1957). For example, before one adapts to a newly filled

eyeglass prescription, one might experience a brief period when rectangles *look* trapezoidal; straight lines *appear* slightly bowed; objects in the periphery *seem* to move when one turns one's head, distant objects *appear* distorted, and the floor *feels as if* it undulates when one walks down the hall. It can make one seasick. Thankfully, adaptation proceeds quickly, and soon one can abandon all these verbs of appearance. Thereafter, the floor feels, and is, motionless and flat; rectangles look, and are, rectangular; and so on.

To use C. D. Broad's (1927) phrase, phenomenal properties characterize the "facts of sensible appearance": how things (in the broadest sense of the word) appear to the senses. One can attribute such a phenomenal property without committing oneself on the question of whether its target is as it appears. Something that (merely) looks blue can look exactly like something that is blue. Judging from how it looks, a witness unaware of any peculiarities of the circumstances of perception would judge the thing to *be* blue. But the appearance is not merely a mistaken judgement; even after one learns of the conditions that make this non-blue thing (in these conditions) look blue, it will for all that continue to *look* blue.

This gives us a clue for unpacking the traditional notion of phenomenal properties (see Clark, forthcoming). There is a similarity in appearance between this thing, seen in these particular circumstances, and some other, avowedly blue thing, seen in other circumstances. At the limit this similarity yields a match, or even phenomenal identity. If two things appear the same (in some respect), then they share a property of appearance. Episodes of perceiving those things present the same phenomenal property. For example, if stimulus x looks exactly the same as stimulus y , then x and y share visual phenomenal properties. If y feels exactly the same as z , then y and z share tactile phenomenal properties.

The entities to which the properties of appearance are ascribed are, paradigmatically, things one perceives (or more broadly, portions of the perceived world), but a person's use of a verb of appearance, instead of the standard copula, flags some oddity in the circumstances of perception, which might be responsible for some aspects of that appearance, and which are such as to call for caution in what would otherwise be a straightforward perceptual judgement. Two portions of the world, presented under different circumstances, might present the same appearance, but in one of them that appearance is more a function of the circumstances of perception than of the thing perceived. It looks funny because of this light, or it appears oddly shaped because of my new eyeglass prescription. Change the light or the prescription and the appearance changes as well. A speaker's awareness of such circumstances is marked by the choice of verb.

The limit of sensed similarity is absolute indiscernibility, but many variants of similarity relations are pressed into service: matching, discriminability, discernibility, relative similarity, recollection of similarity, and all their negations and contraries. Traditional phenomenal properties are individuated by appeal to the discriminations that a subject can make among the conditions

that it perceives. Stimuli that are absolutely indiscernible present the same appearance; stimuli that can be discriminated from one another present at least slightly different ones. In this way a traditional phenomenal property--an attribute of sensible appearance--is tied to capacities to perceive similarities and differences among stimuli.

Capacities to discriminate can provide tests sufficient to attribute or deny a (traditional) phenomenal property. If subject *S* can routinely and reliably discriminate between occasions in classes *P* and *Q*, then stimuli in class *P* do not present the same appearance as those in *Q*--and *S* can perceive the difference. If for example she can reliably discriminate line orientations that differ by just one tenth of a degree, then to her the orientation of one line in such a pair must not appear the same as that of the other. If she can reliably discriminate spectral light of 580 nm from light of 584 nm, then however they look we know those two wavelength packets do not look the same to her. A test for whether some human females have a four dimensional color space is had by testing whether they can reliably discriminate differences (and relative similarities among) a class of stimuli which to normal trichromats are all mutually indiscriminable.

One way to summarize what seems paradoxical about blindsight is: subjects can pass many of the tests that would typically be taken to show that they are perceiving something, and that they can discriminate it from other stimuli, even though they sincerely deny being aware of any aspect of the stimulus in question. Consider the evidence mentioned in section I. DB can with surprising reliability discriminate location, X's from O's, and the direction of motion, even though the stimuli are presented in such a manner that he is not aware of any feature of them. The subjects in the Stoerig and Cowey experiments could with surprising reliability "guess" whether the light was "yellow" or "orange" even though they were not aware when the stimulus was present, and had to be prompted when to "guess". In fact, I will argue, they pass some of the tests that indicate they register the chromatic appearances of stimuli within the scotoma. An alternative phrase: they pass some tests that indicate they have states with qualitative character. They pass such tests even though they are not aware of the objects (or stimuli) presenting those appearances.

IV Mysteries inside Enigmas

Clearly DB is discriminating some differences between stimuli presented in his scotoma even though (if we take him at his word) he is unaware of those stimuli, and must be prompted to guess. The simplest hypothesis, which I will favor, is that the difference he discriminates between stimuli *is* a phenomenal difference. It is simply one of which the subject is unaware. Otherwise it has all the characteristics of your run-of-the-mill difference in phenomenal properties. To stand the Weiskrantz quote on its head: DB has the very qualia themselves, but is simply not aware of them. Why not?

The question is apt to met with incredulity, sputters of outrage, and then

finally an argument: "But that's not what the word means! Qualia must be conscious. The proposal that there exist unconscious phenomenal differences, or that DB is being appeared-to in a matter of which he is unaware, is speculative, useless metaphysics. It distorts language and can only sow confusion."

I will repeat that one can, if one likes, reserve the words "qualia", "phenomenal character", and "appearance" to just the ones of which someone is conscious. But the critical question is whether or not there can exist a family of other properties, which differ from these "qualia", "phenomenal characters", and "appearances" in *only one* respect: that no one is aware of them. Could there be differences, that are exactly like those differences that you call "phenomenal", excepting only the fact that no one is conscious of them?

The question needs some tightening. Of course there will be many consequences--logical consequences, psychological consequences, and other nomological consequences--of the fact that no one is conscious of these differences. For example, the particular instances will not be reported by the subject in question, or by any subject; they will not be described in English, or in any natural language; they will not appear in that subject's autobiography, if such an autobiography is ever produced. They will inspire neither sonnets nor self recriminations. Some philosophers will urge it follows as well that they are not in any sense phenomenal. Lump all these consequences together and call them "Consequences of Non-Consciousness". Could there be a family of properties P such that *all* the differences between P and the normal, conscious, phenomenal properties Q are found among the Consequences of Non-Consciousness? Could there exist such creatures of darkness? If you think the answer to that question might be yes--that a "yes" is conceivable--then you are open to the possibility of a "non-standard" interpretation of the relations of chromatic qualia to awareness. If you think the answer must be no, then I fear you have erected an a priori roadblock against the possibility of our ever understanding phenomenal consciousness.

Consider the implications of a "no". To say "no" is to say that it is not the case that there exists a family of properties whose differences from the normal qualia (or normal phenomenal properties) are all found among the Consequences of Non-Consciousness. So for any quale Q , any such candidate property P must differ from Q in some other respect besides those consequent upon the fact that someone is conscious of Q , and no one is conscious of P . In other words, quite apart from the fact that Q characterizes something of which someone is conscious, and P does not, and apart from all the consequences of that fact, there is also some *other* difference, or differences, between them.

This implies in turn that that of which one is aware when one is aware of a quale cannot be identical to any property that can exist even if no one is aware of it. All the candidates must differ, not only in not being conscious, but also in some other respect besides those consequent upon non-consciousness. One might think that one could give a characterization of that of which one is

conscious that is independent of the fact that one is conscious of it. In the typical two-term relation, one can name the terms independently of their relationship to one another. But the claim that qualia *must* be conscious implies that here that rule fails: qualia cannot be identified with any properties for which such an independent characterization is possible. Not only is the consciousness of them a mystery, but that of which one is conscious when one is conscious of them is a mystery as well. A solution to either requires the other already be solved. This is not a mystery inside an enigma, but something worse: two mysteries, each of which is *wholly* inside the other.

In contrast, a "yes" answer allows that there might be candidate properties that differ from qualia in (essentially) just one respect: that no one is conscious of them. The very same difference of which a subject is aware when a subject is aware of the difference between the appearance of yellow and the appearance of orange can exist even if no subject is aware of it. If so, it is possible--conceivable--to analyze a phenomenal property of which a subject is conscious independently of coming to understand what it is to be conscious of it. This mimics a strategy applied to understanding many perceptual capacities, such as what it is to see a color. Eventually, after many centuries, our kind came to understand some of the physical characteristics of the stimulus--specifically, of the electromagnetic radiation impinging on the retina. That understanding at last allowed investigators to explore how a physical nervous system might detect and discriminate variations in those physical characteristics. It unlocked the doors. Without that characterization of the stimulus for color vision, it is hard to see how human color scientists ever would or could have made what progress they have made in understanding human color vision.

A similar divide-and-conquer approach applied to awareness of phenomenal similarities and differences suggests that we first try to understand that of which one is aware, independently of the awareness of it; and then try to understand what makes it true that some subject *S* is aware of some such thing. Blindsight, hemi-neglect, and other disconnection syndromes strongly suggest that differences in sensible appearance can be registered in various pre-conscious, preattentive perceptual processing mechanisms. This registration is not yet "phenomenal consciousness" but perhaps all that is missing is awareness and what follows therefrom. If so, then we could characterize that of which one is aware independently of the awareness of it. To such registrations all we need add are capacities that suffice sometimes to make one conscious of what is registered.

Can this latter step ever be completed? The audience (including me) at last confronts the question (Larry Hardin's question) squarely in the face: what must we add to pre-attentive perceptual information processing in order to instantiate a subject who is aware of a quale of color?

V Selective Attention

Surprisingly enough, I think we can start to provide parts of an answer to this

question. These parts derive from a philosophically neglected psychological capacity: the capacity for selective attention. How does contemporary cognitive psychology treat the phenomena of selective attention? The first surprise (for philosophers) is that the emphasis is very much on the word "selective". It is thought that there exist many independent, peripheral, pre-attentive perceptual processes that operate in a parallel, bottom-up fashion, for many different modalities, and different channels within modalities, simultaneously. It is also thought that there exists relatively more central, lower-bandwidth, general-purpose processing, which can absorb inputs from all the different peripheral processes, and integrate them to form one action-plan for the organism as a whole. Somewhere between peripheral and central processes, it is thought that there must be some process that selects which of the peripheral processes are going to receive the benefits of further, central processing. Psychologists fiercely debate the details of the characterization of peripheral and central processes. They debate where, exactly, the selection process occurs. They debate why the selection occurs. But it is (almost) universally acknowledged that there is some selective process, somewhere. That selection is the job of selective attention.

One old and somewhat discredited explanation for why selection is required appeals to the sheer differences in bandwidth--in bits per second--of all the perceptual data coming in per second, compared to what a central processor might be able to manage. During every waking moment all the sensory modalities are firing away all the time, transmitting impulses up their many tracts into the central nervous system. These transmissions are initiated peripherally, as independent and "bottom up" processes, and while a certain amount of central (and peripheral) gating is possible, there is no central way to turn off receptors. The resulting volume, even in a single modality, beggars the imagination. For example the optic nerve in humans contains about a million fibers for each eye, and each nerve can (theoretically) fire up to a thousand times per second. If we simply want to record for each millisecond whether the nerve has fired or not, we require a recording capability for each optic nerve of a gigabit per second.² And what does one do with all the data? The typical scholarly monograph requires roughly a million typed characters (a megabyte), which with endnotes and references takes roughly ten megabits to store; and so the optic nerve from one eye could theoretically transmit a small library, of one hundred such volumes, each second. Over the years book storage becomes problematic even at very low annual rates of acquisition, and here we are adding a hundred per second.

And, remember, these are transmissions from just one eye; we have the

2 The problem with this sort of calculation is that no one knows what the relevant ensembles are at the two ends of the channel: how the information might be coded or "chunked". It is almost certainly not the case that the lateral geniculate nucleus needs to record data at the rate of one thousand bits per second for each nerve in the optic tract. Probably visual information is coded and chunked in ways that does not require such high fidelity. But, as mentioned later, auditory information may in fact require such high fidelity, since minute differences in the timing of wave crests arriving at the two ears are used to determine the direction of the sound.

other eye, and all the other senses, to manage as well. Humans have about a million muscle spindles, each with its own neural control. The auditory nerve may contain less than a hundred thousand neurons, but it almost certainly does use the full temporal resolution that is possible, so even a mere hundred thousand neurons might yield a bandwidth of a hundred megabits per second. No one quite knows how to assess the information-bearing capabilities of dynamic touch and the somesthetic modalities, much less taste or smell. But in any case it is clear that if one considers the transmission lines in a pure information-theoretic way, their sum remains many orders of magnitude beyond the capacities of any computer available today.

The necessity, then, is to select. It is hard enough to keep up with the current rates of academic book publication; it is humanly impossible to keep up with a hundred books per second. The trick lies in selecting what to ignore; to pick the relevant bits out of the stream, and focus one's efforts on them. This selection is the key job of what cognitive psychologists think of as "selective attention". One set of analogies likens it to a gateway or filter, through which only selected streams of peripheral information are allowed to pass.³ The sources of these streams are conceived as parallel and autonomous sensory processes, which mostly proceed in data-driven, independent, and bottom-up fashion. The destination for this information is "central processing". It is "central" in various senses: it is not scattered about in multiple, independent, parallel processing modules; it is a common destination for all perceptual representation; it makes its results available to all portions that are "central"; and, relative to the amount of processing power that would be needed fully to analyze all that perceptual information, it has a limited capacity. This limited capacity may give it the appearance of proceeding serially, one step at a time, rather than in multiple, parallel streams. Furthermore, it is "central" in that its resources can be allocated and reallocated under some degree of central control; more or less all of them can be assigned to or focused on one or another task. The function of selective attention is to select which information is to be allowed to pass through to this further, central, processing.

This picture of the landscape lying between perception and attention has undergone elaborations and amendments since it was first proposed over fifty years ago (see Broadbent 1958, Neisser 1967, Fodor 1986), but in broad outline it is a well established part of the conventional wisdom of cognitive psychology. In particular, the need for *some* sort of selection *somewhere* between peripheral sensory processes and central cognitive ones is almost universally acknowledged, and that selection is acknowledged to be the job of selective attention. This job specification can help us to understand some of the powers that folk psychology ascribes to attention. For example, it is a truism that if you focus your attention on something you perceive, you can

3 Other analogical models include gain control, amplifier, zoom lens, telephoto lens, and channel. Some of these analogies emphasize amplification of the selected signal; others the inhibition of the non-selected signals. See Laberge 1995, 39.

often perceive more of its details, perceive it more quickly, and perceive it more reliably. A good reason to advise novice drivers to pay attention to the traffic is that doing so helps them to anticipate and respond more quickly, and thereby avoid some accidents. Cognitive science provides an explanation: by "paying attention" to the traffic the novice driver opens certain switches, gates, and channels within the nervous system, and inhibits others, so that visual information (and not other kinds) gets the benefit of central processing. This further processing allows finer details to be resolved within the flow of visual data, and allows for a quicker, more accurate, and centrally coordinated response to exigencies that are detected visually. In folk terms: if you pay attention to what you're looking at, you can see things better.

We can also demystify some of the persistent optical metaphors applied to selective attention. It is often compared to a "spotlight": the stimuli to which one attends are "lit up" by a spotlight, which is under some degree of voluntary control. Similarly, consciousness is often described as an inner light, which lights up certain processes, leaving others in the dark (see Rey 1997, 472). The light picks certain items out from the background, and enables one to see them better. Now it is true that more illumination can help one see better, but additional processing power devoted to the task can do the same thing. As anyone who has used a digital camera knows, the more processing power one can throw at a picture, the quicker its details can be resolved. So there is more than one way to "see" better. The "spotlight" analogy is probably as old as the human capacity to direct artificial illumination upon things, but wider use of computers may someday provide it with a rival. "Focus" itself provides a second optical analogy: only when it is held at the correct distance will the magnifying glass bring the object of interest into focus, but when it does the view provides more detail and finer structure. Again the analogy points to a way of receiving more information from the object of interest, but better processing of the same information could yield the same effect.

VI. The act/object divide

Psychologists, then, emphasize the selective functions of selective attention. The reason this is something of a surprise to philosophers is that philosophers tend to pay attention to other aspects of attention: in particular, to its intentional and quasi-perceptual attributes. In ordinary language attention is described as picking out an "object", in something like the way an intentional state represents or is about an intentional object. We speak of focusing attention on an object, or directing attention at an object. The object is picked out by the attentional state in something like the way an object one perceives is picked out by a perceptual state. Such an object need not be a physical object or a "thing" in any ordinary sense of the word "thing". It should be read as an intentional object: as whatever that state represents or is about. So the "object" to which one pays attention might be a physical thing, but it might also be a feature or property, a set of relations among features, a state of affairs, a region of space, an event, a process, or whatever. This intentional or quasi-perceptual

aspect of attention is what grabs the spotlight in philosophical discussions.

Some theorists think of attention as functioning something like a demonstrative: it is a way of picking out some of the objects one perceives without requiring a description true of just those objects (see Pylyshyn 2001, Campbell 2002). This "picking out" is done directly, in something like the way the "this" is assigned a referent when a speaker makes a successful demonstrative use of the word. The referent must be among the objects currently perceived, but that condition does not suffice, since many things are perceived to which it does not refer. Typically there is an additional demonstrative gesture that helps to pick out the referent. Likewise, many objects might be currently perceived, but attention somehow picks out from among them some proper subset, whose members comprise the current "focus" of attention. The optical analogies give a perceptual model for this picking out: more light is shone upon the objects to which one attends, or they are in better focus. In this way the objects to which one attends would be perceptibly different from ones to which one does not attend, and the difference between attending and not attending seems to be a property of how they appear. They are "lit up" or "in focus".

Now if the psychological consensus described in the previous section is correct, then the real nature of selective attention--its scientifically discovered real essence--is that certain channels are opened or enhanced, and others are closed or inhibited, so as to allow a select subset of sensory representations to enter central processing. To pay attention to an object x is to open certain gates and switches in the nervous system so as to allow representations of x to be further processed in "central" portions of one's cognitive architecture.

This explains an interesting feature of the locutions we use when we talk about attention. Those locutions seem to speak of objects, but when one talks of "attending to an object" one is certainly not attributing any kind of property to that object. Instead one is talking about relations among psychological states that are directed at, or that are about, that object. For example, when we say "Sam is paying attention to the fishing line" we are not saying that some ghostly illumination shines on the fishing line, or that it shines somewhere inside Sam's head; instead ostensible reference to the fishing line is used to characterize relations among some of Sam's perceptual representations. In particular: representations of that object are the ones selected to receive the benefits of Sam's central processing. So Sam is more likely to react quickly if the line twitches.

Our ordinary locutions for "attention" employ this handy shorthand. We mention the "object" of attention as a way of picking out a set of representations whose intentional object is the one mentioned. So to "pay attention to" x is to alter the configuration of gates and channels inside one's head so that representations of x , and not of other things, receive the benefit of further central processing. That central processing in turn allows one to resolve more sensory and structural detail in x , perceive it more quickly, remember it better, and so on. But how one actually manages to open the right

gates and channels, and close or inhibit others, is really something of a mystery. By "paying attention" to object x one must succeed in opening particular gates and switches in one's head so as to allow those representations that are representations of x to pass through to central processing. How one controls such gates and switches (or for that matter, how one *finds* the appropriate ones to control) is still something of a mystery, but that's what the exercise of selective attention *does*.

VII. Perceptual Awareness Demystified

The fact that selection by selective attention is competitive imposes some significant structure upon what might otherwise be a gaggle, a horde, an unruly mob of on-going mental processes. One implication, for example, is that some of those processes are *not* selected. Put in folk terms, at any given moment you may be paying attention to something you perceive--the twitches on the fishing line, for example--but there are certainly plenty of other things, features, processes, or events you perceive at the same time to which you are not paying attention: the sound of your breathing, the pull of gravity, the taste in your mouth, and so on. One feels the pull of gravity, and one could focus attention on it, but attention is focused elsewhere. It would be handy to have a label for all these on-going sensory processes which could have been selected by selective attention, and would have been, but for the selection of something else. I suggest that we *do* have a label for sensory processes that have exactly this status: we say that those are phenomena of which we are *aware*. These are processes that are in the running for selection, but happen not to be first at the moment.

The idea is that selective attention picks out a subset of all the streams of perceptual processes currently underway, and the objects of those streams are the ones in the "spotlight" or in the "focus" of attention. But one also perceives objects not currently in the spotlight. One must be able to gather some information about such objects, as otherwise attention could never shift away from its current focus in an adaptive fashion. One sense of the ordinary term "aware" applies precisely to all those objects not in the spotlight, but in the penumbra: all the ones perceived, and potentially selectable, but currently not winning the competition for selective attention. You are "aware" of them even though you are not "attending" to them.⁴ Often the only thing precluding such a win was the fact that something else won instead.

There are various bits of evidence for this account of the distinction between "attending to" and "aware of". Attention is directed at some "focus", some selected subset of the objects currently perceived; and the directing of it is under some degree of voluntary control. Otherwise a teacher who raps the table and says to the class, "Pay attention!" would be issuing an imperative that

4 Or at least not *focally* attending to them. In the end, probably both sets of phenomena are products of "divided" attention: of the allocation of attentional resources. The ones in the "spotlight" win the plurality of those resources, while those in the "penumbra" get a tiny share. See Dehaene & Naccache (2001).

is absurd. Indeed the class may in time learn how to better focus its attention, though such outcomes seem rare these days. Whereas "being aware of" applies indifferently to a much larger field; it is not directed at anything in particular. One cannot "focus" one's awareness, unless this just means focusing one's attention. The spotlight metaphor gets no grip on our locutions for awareness: phenomenologically we do not find a foreground of things of which one is aware surrounded by a background of things of which one is unaware. Phenomenologically, such a background is equivalent to no background at all. Nor is awareness under the same degree of voluntary control as selective attention; a teacher who raps the table and says "Be aware!" would be issuing either an absurd imperative or a joke (unless it means simply "Beware!" or "Pay attention!").

The suggestion is that "aware of" applies to all the objects of representations that are in the running for selection by selective attention but which are not the current winners.⁵ This can be a large field, as it includes all of the also-rans. But one does have to be in the running. Precisely because it includes all those that are *not* winning, it has no foreground/background organization, it has no particular focus or direction, and so it cannot be directed deliberately at anything. There is other evidence for this analysis. Etymologically, the word "aware" comes from the same root as "wary", "warn", and "beware", denoting the exercise of at least some caution, care, preparedness, wariness, vigilance, or responsiveness. (OED: [OE *wær*, *gewær* corresp. to OS *war*, OHG *giwar* (G *gewahr*), f. Gmc base meaning 'observe, take care'. Cf. *WARE* *n.*2, *n.*3, *WARE* *v.*1, *WARN* *v.*1]) One is aware of all the things that could have attracted one's attention in the prior moment or might do so in the next. A tripwire is in place, which if tripped would attract attention, which in turn could call up concerted action of some sort--a sort to be determined by central cognitive resources. In contrast, a guard caught unawares is exercising no such wariness; the bad guys can sneak in without the guard noticing, without catching the guard's attention, without alerting the guard. I suggest that the level of vigilance described by the former locutions is: events *could* attract selective attention; the only impediment to their success in doing so is that something else is, at the moment, dominating the selection. The tripwire in place (if one is "aware" of one's surroundings) is one that simply draws attention; it does not lead directly to any actions except those of noticing, alerting, and attending-to.

Whereas if one proceeds unawares, then even those actions do not occur. In the guard caught unawares the channels between the receipt of sensory information and the engagement of selective attention, are, unfortunately, closed, or at least attenuated.⁶ To say "the guard was not aware of the sound"

5 Again, recent models suggest that such selection is not all or none; the "winner" just wins a plurality of the available attentional activation, while those that remain "in the running" receive just enough to keep them active.

6 Since selection is not all or none, one should not think of representations passing through the selective attention "gate" on the analogy of a cow passing through a gate in the fence. The gate might simply weaken the signal strength of the cow. (See the other analogies mentioned in note 3.)

implies that the guard's auditory representation of the sound was one that could not draw attention, and so was not a candidate for entry into central processing. If one is completely unaware of the sound, one is incapable of shifting attention to it. The cost is that no central planning can be aroused.

Now according to the argument in the previous section, our talk about the "objects" of attention is a way of referring, indirectly, to a status enjoyed by representations *of* that object: that those representations are selected for further, central, processing. If this is so then our talk about the objects of which we are aware is likewise a kind of shorthand. To say "Sam is aware of the slosh of water in his boots" implies that at any moment Sam *might* focus his attention on the water sloshing in his boots. He could do so. But mention of the putative object is a means of indirect reference; the ultimate translation is of the form "Sam's somesthetic sensations of water-sloshing-in-boots are in the competition, and could be selected for further, central, processing; but at the moment are not in first place." Talk about the "awareness" of something is a way of talking about the status of perceptual representations of that thing; in particular, their status in the competition for selective attention.

An analogy from economics might help. Such talk characterizes what might be called the informational "liquidity" of perceptual representations. If the representations have high liquidity, the information therein is transmissible; conveyable; mutable; it can be deployed and employed elsewhere. Channels leading from the representations are open and ready and could, with but a twitch in the direction of attention, send them downstream to engage central capacities, alter responsiveness, and so on. One can speak of the downstream subpersonal modules as "consumers", except one must keep in mind that information, like capital, is rarely "consumed". It is not so much "consumed" as distributed, manipulated, analyzed, invested, stored, re-invested, saved, and recycled again. Here and there it is wasted, squandered, or lost. So we have not just downstream "consumers" but investors, analysts, loan agents, manufacturers, profiteers, distributors, and so on. Some minds contain wise guys and loan sharks. In terms of this analogy, awareness does not alter the value or even the currency of the capital received; instead it changes the routes through which that same capital can be deployed. It changes the liquidity of that capital: its potentialities for deployment, investment, re-investment, and so on. Some types of capital are more liquid than others; and given a fixed type, different economic arrangements can make that same type more or less liquid. One implication of perceptual awareness is that some perceptual representations have high liquidity. They are perched in such a place that they could, with just a flick of the switch, be transmitted and invested elsewhere. That is, dropping the analogy, they could engage selective attention.

VIII Preattentive registration

The winners of the competition to attract selective attention are the ones we attend, scrutinize, notice, or heed. We are "aware" of the ones that are

competing but not winning. But notice there is a third category. What of all the ones who do not even enter the race? These would be sensory representations to which one could not shift selective attention; for one reason or another they are not even competitors. What of them? What is their status?

I will argue that this category is indeed a robust one, and it is filled with states that occur during any episode in which perceptual information about something is preattentively registered even though the subject is not aware of the thing in question. Various circumstances can produce such episodes, such as those found in masking studies, priming studies, search under extreme time pressure or high cognitive load, covert recognition, and perhaps hypnotic suggestion. And this at last is the category that can help us to understand what is happening in the disconnection syndromes of hemi-neglect and blindsight. What we need for them is precisely a category of what might be called "mere" sensing: specifically, of sensory registration without awareness.⁷

Such a category is, first, logically possible. To revert to an earlier example. Suppose we find some specific set of conditions under which it would be impossible (it would contradict the laws of psychology) for Sam (or any human in the same situation) to focus attention on the feeling of water sloshing into boots. Selective attention has operating parameters, and they can be exceeded, jammed, or abused. Selection is not purely random, and is based in part on information about current goings-on. The mechanisms that determine where attention will go next must also have access to such information. (If selection is based on location, the mechanisms driving selective attention must be able to pick out where the feature in question is located. If selection is object-based, there must be information available to identify the relevant object.) In any case the selecting itself requires current information. Now suppose that our conditions are carefully designed to overflow the information buffers, or cause some other glitch in the operations, so that at that moment the selecting process literally cannot select the representations of water sloshing into boots. Our conditions are somewhat similar to finding a "bug" in the operating system.⁸ Perhaps a task overload or "denial of service" attack will suffice: Sam is simultaneously landing a

7 Psychologists these days are rightly leery of the words "sensation" and "sense", as they suggest to some ears an appeal to an old and discredited theory according to which each perceptual modality starts with elemental states, accessible to consciousness, called "sensations", and from them builds up "perceptions". I do not mean to endorse that theory by using the verb "sense". Nor do I mean to imply that those processes are either elemental or accessible to consciousness. The verb has not yet been banned from philosophical discourse, and there I think it sometimes means something close to what is suggested here: the preattentive registration of perceptual information. Note that showing sensation without awareness is something less than showing *perception* without awareness; in particular we are accustomed to the idea that information extracted in early sensory systems is not invariably available for the guidance and control of action. However, many of the problems that Dretske (2006) found in experimental tests for the *lack* of awareness infect the sensory domain as well.

8 The so called "attentional blink" provides a real example of pushing the envelope of attentional operating parameters. There is an attentional "dwell time" of several hundred milliseconds (which may reflect the time it takes to disengage and move attention); if a second stimulus which needs attention is presented during that interval, error rates go up enormously. See Raymond, Shapiro & Arnell 1992, 1995.

difficult fish and trying to complete a set of nearly impossible cognitive tasks designed by a clever cognitive psychologist. With tasks sufficiently well tuned to exploit particular weaknesses in the mechanisms of selection, the sensing of wet feet might get knocked right out the race. Perhaps, given that task load, it is not even in the running, and no normal human in that situation could at that moment shift attention to feeling the slosh.

If in this sense Sam *cannot* turn his attention to the water sloshing into his boots, then common sense would suggest he is as good as insensible to that state of affairs. He certainly cannot notice it, or observe it deliberately, or scrutinize it, since all those require success at attracting attention. For the same reason the event is not one that Sam could report or describe; nor could he use a demonstrative to refer to it. He could not use "*That* water", even in thought, to secure reference to the water in question. He could not satisfy any of these ordinary indices of being *conscious of* water sloshing into his boots. His overt actions would not differ from those of someone who does not feel the water at all. In fact he could not do anything deliberately--could not perform any voluntary intentional action--that would indicate that he had felt the influx.⁹

Common sense makes no distinction between Sam being insensible to, and Sam being unaware of, the water sloshing into his boots. There is indeed a sense of "insensible" which is, simply, "unconscious". It is not surprising that these two concepts are often conflated. Clearly enough one cannot be aware of a sensible feature (like the feeling of a slosh) unless one senses it. But common sense perforce relies on overt behavior and verbal report; the slow progress of experimental science in this domain has also gradually provided covert and indirect indicators, which can tease apart the two notions. So for example skin conductance can indicate "covert" recognition in someone whose prosopagnosia eliminates all overt or verbal indicators of recognition (such as naming the person.) Similarly, one can get indirect evidence (or covert measures) indicating that Sam is not strictly insensible to the water (he does indeed *sense* it: the information registers somesthetically), but simply that he cannot focus his attention on it. He "senses" it perfectly well--the preattentive registration of somesthetic information proceeds normally--but here there is a hitch in the attentional mechanisms. Such a case would be one in which he "senses" the water but is not aware of it.

If you are persuaded that this state of affairs is possible, one more step will complete my argument. Within that envisioned state of affairs occur psychological states of preattentive registration of perceptual information. Suppose that given normal background conditions, some of those states are necessary and sufficient to subserve the discriminability of, and relative similarities among, various sensible features. Discriminations hang on the presence or absence of these particular mental states. For example, they carry

9 Such actions are precluded because they require attention to the somesthetic event. The impossibility in question would also rule out an "exogenous" shift of attention--one in which the subject does not deliberately shift attention, but instead the subject's attention is drawn to some stimulus (for example, by movement or change).

the information sufficient to render the feeling of water sloshing into boots discriminable from the feeling of (say) mud, or motor oil, or apple sauce, or some other liquid sloshing into boots; and likewise they carry information necessary and sufficient to discriminate water sloshing into boots from water sloshing into trousers, or sleeves, or hat. Now, as argued above (in section IV) if we are ever to get a non-mysterious account of phenomenal consciousness, at some point we will find a family of states and properties that differ in (essentially) just one way from the normal phenomenal properties of which one is aware, that difference being: that no one is aware of them (and all that follows from that). Suppose in the fullness of time we identify the preattentive states that subserve the discriminabilities and relative similarities of the sensible features in question. We would then have grounds to conclude that those states, even though they are preattentive, register *phenomenal* differences; all they lack is some standing in the competition for selective attention, which (per hypothesis) would make one aware of what they represent.

A terminological variant on this conclusion is that these states of preattentive registration of perceptual information are states with qualitative character.¹⁰ They have all the properties of the normal ones, except those that follow from the fact that no one is aware of their objects. DB might have "the very qualia themselves", although he is not conscious of them. To use Chisholm's terminology: DB is in a state of being "appeared-to", even though he is not at all aware of the appearances.

The suggestion is that states of being appeared-to, states that have a qualitative character, need not be states that make someone conscious of something. The concepts "being appeared-to" and "being conscious of" are distinct and independent. If so, something can look red to someone even though that someone is not aware of the thing that looks red, of the reddish look of it, or of seeing something that looks red. The only sense in which such states are "states of consciousness" is simply that a creature who has them is not an entirely unconscious creature, since some mental activity is going on.¹¹

10 I follow the terminological convention of treating "qualitative" properties as properties of mental states, and "phenomenal properties" as characteristics of sensible appearance, which are attributed rightly or wrongly to the things one perceives. So how the water feels when it sloshes into one's boots is a phenomenal property, characterizing how a portion of the world appears. Whatever properties of somesthetic sensory states subserve the discriminability of that feeling from others constitute the qualitative character of those sensory states. Unfortunately the term "phenomenal character" is often used in the literature as a synonym for "qualitative character"; I use it that way myself at the beginning of this paper.

11 One might say that these states have "phenomenality" even if they are not cognitively accessible (see Block 2002, 394). (Read "phenomenality" here as: qualitative character sufficient to make something one perceives appear a certain way.) Or one might say that their "phenomenology overflows accessibility" (Block 2007, 487). But the terminology is treacherous. I do say that bottom-up preattentive perceptual systems have a far greater bandwidth than central processing (section V), and that some states within those preattentive systems suffice to sustain the presentation of phenomenal properties (section VIII). These yield the conclusion that some systems sufficient to sustain the presentation of phenomenal properties overflow (have greater bandwidth than) that which is cognitively accessible (e.g. that which can be processed in the global workspace). In places Block's conclusion is put in convivial terms: "If we assume that the strong but still losing coalitions in the back of the head are the neural basis of phenomenal states (so long

Likewise, states of "being appeared-to" are in that minimal sense states of consciousness: they indicate that the animal in question is not comatose or asleep. But almost any mental activity suffices to show that.

A likely response to this suggestion is: but how can one be appeared-to without being conscious of the appearances? I admit the suggestion is weird, but that is all it is: weird, uncommon, improbable. By hypothesis your own states of unconscious sensation are not accessible to you. The proposal is logically analogous to Freud's proposal (1900) that unconscious wishes exist: it is not refuted by angry members of the audience declaring that they are aware of no such wishes in themselves. Freud delighted in presenting the hypothesis to the public, and took the many expressions of anger as evidence *for* his hypothesis, though this latter inference was not sound. One can simply turn the question on its head: why must every appearance be an appearance of which someone is conscious? Being appeared-to and being conscious-of are distinct concepts; what is the necessary link between them?

IX Two Objections

One response to this question might be put as follows: if these really are episodes of being-appeared-to, instances of mental states that have some qualitative character, then we ought to be able to answer the question: *who* is being appeared-to? How can an appearance be "presented" if there is no one *to whom* it is presented? As noted above (section II), Nagel himself argued that phenomenological properties, like consciousness, are subjective: they can only be understood from the "point of view" of the subject who has them (Nagel 1974). So perhaps what links "being appeared-to" and "being conscious of" is that both of them are subjective phenomena. They can only be understood if one understands the "point of view" of the subject who has them.

Now there is a sense in which it is quite clearly true that in order to understand phenomenal properties, we have to understand the point of view of the creature whose perceptions instantiate those properties. Namely, we have to understand the similarities and differences which *that creature* senses to obtain among the stimuli that confront it. These are not determined by the stimuli themselves. As earlier emphasized, the similarities and differences of

as they involve recurrent activity), then we have a neural mechanism which explains why phenomenology has a higher capacity than the global workspace." (Block 2007, 498). But there is at least one disagreement: I see no reason to believe that those preattentive states which have qualitative character are necessarily "experiential" or that the subject is in any sense "aware" (or "Aware") of them. Even though Block is now "abandoning the term 'phenomenal consciousness'" (2007, 485), for him states that have "phenomenology" still seem to be "experiential" states. He describes subjects in the Sperling experiment as having "experiences as of specific alphanumeric shapes" (2007, 287 (section 9) and says "The argument of this paper depends on the claim that subjects in the Sperling and Landman experiments have phenomenal experiences of all or almost all of the shapes in the presented array." (Block 2007, 491, section 11). Finally, activation of the fusiform face area in patient G.K. is described as possibly yielding "phenomenal experience that the subject not only does not know about but in these circumstances cannot know about." (Block 2007, 498, section 14). If these three citations were amended, and the assumption that qualitative states are necessarily experiential were dropped, I think there would be very little difference between Block's conclusion and the one here.

phenomenal properties need not and do not map in any direct way onto similarities and differences of the physical properties of stimuli. Instead one has to understand how the creature's sensory systems work, so that one can understand the similarities and differences *it* perceives among those stimuli. Part of the objection is hence well founded: we do, in a certain sense, need to enter into the "point of view" of the creature who has these sensations.

But the sense in which we need to do this does not sustain the rest of the objection. In particular this "perspectival" character of phenomenal appearance is not sufficient to yield "subjective facts", or facts of a different ontological or metaphysical kind than ordinary facts. Nor does it license the inference to the conclusion that there is no appearance unless there is someone or something *to whom* the appearance is "presented". The "presentation" terminology was initially introduced as a handy way of stating the converse of the relation "being acquainted with", and so it does sound odd to say that DB is "presented the appearance" of something even though he roundly denies being aware of any such thing. It sounds less odd to say he simply "has" a sensory state that underwrites the successful discrimination. That state stands in the same place in the sensing of similarities and differences--it has the same qualitative character as--the states that normally makes one conscious of seeing those things. But in DB that very state does not make anyone conscious of anything.

A second objection might be put as follows. Clearly states of preattentive registration of perceptual information are not states of phenomenal consciousness. I have urged that adding the capacities of selective attention yields an architecture that is closer to one to which we are willing to ascribe states of phenomenal consciousness. But how do the additions do that, exactly? Let us imagine this in some detail. We might construct a special purpose wine-tasting machine, for example. One pours a sample of some vintage into the input funnel. The machine takes various measurements, computes some numbers, and eventually types out "a chemical assay, along with commentary: 'a flamboyant and velvety Pinot, though lacking in stamina'--or words to such effect" (Dennett, 1988, 46). This machine might have oenological discriminative powers that far exceed any human's, but at best it has the functional equivalent of preattentive sensory registration. It is not aware of anything it "tastes", nor is there anything it is like to *be* that wine-tasting machine. All the processing proceeds in the dark.

Now to this machine we add some additional circuitry, which adds the functional equivalent of the mechanisms of competition and selection, and of the releasing, shifting, and maintaining of attention on some selected subset of the representations coursing through its innards. How does any such addition make it any more reasonable to say that the machine is now aware of what it "tastes"? Couldn't all the latter proceed in the dark as well?

There is room for but the barest sketch of an answer. Notice first that if a psychological model of perceptual experience is thought to be possible at all, then if our search for it goes well, at some point it will be reasonable to ascribe such mental states to an instantiation of such a model. After all, per

hypothesis, *we* are an instantiation of such a model. The real question is what one must put into such a model to make it a likely candidate for satisfying such ascriptions. The ones found early in the search are absurdly simple, but that does not show that all the later ones will be as well.

Second, it would be no mean feat to add to any artifact the functional equivalent of the mechanisms of competition and selection, and of the releasing, shifting and maintaining of attention. This would add much more complexity than one might think. For example, to have competition, there must be competitors, implying a range of possible alternative activities to which the relatively limited resources of something that it is fair to call "central" processing could be applied. So a single-purpose, one-track device cannot qualify. There must be a plethora of distinct representational streams to which one might "pay attention". This in turn requires a genuine architectural distinction between peripheral and central processing; a multitude of the former processes must compete for the benefits of the latter. The system must also be able to select one such stream, shift resources to it, keep others in abeyance, yet allow itself to respond to emergencies, or shift to others as needs or priorities change. A system that can "pay attention" must also be distractable; and to say a system is "distractable" is to say something rather complicated about it. It has a finite attention span. Even during that span competitors beckon. It can be heedless or focused. It can become enthralled, absorbed, or bored.

If one did duplicate the architecture of selective attention, then competition and selection, the allocation of central processing, the runner up status of inputs not selected, the monitoring of those runner-ups, and all the rest would be duplicated as well. And I suggest that it is then not quite so easy to imagine such a device as totally lacking awareness. Awareness becomes a live question only if our wine-tasting machine could become distracted or bored, or if it took some effort to get it to focus on the task we would like it to complete--but then, I think, it really *is* a live question. So oddly enough it is the (temporary) absence of focused attention that betokens the presence of awareness. The short attention span of some undergraduates is regrettable, but not because it indicates lack of consciousness. In fact it indicates precisely the opposite: the undergraduate is attending to something *else*. Even an attention span of zero fails to indicate lack of consciousness. If the idioms of attentiveness apply at all, then so do the idioms of awareness. To get to the pure repose of no consciousness, one must entirely eliminate the applicability of our idioms of attention: not an attention span of zero, but no attention span at all.

References

- Block, Ned (2002). The harder problem of consciousness. *Journal of Philosophy* 99(8): 391-425.
- Block, Ned (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences* 30 (5/6): 481-548.
- Block, Ned, Flanagan, Owen, and Güzeldere, Güven (1997) (eds.) *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press.
- Broad, C. D. (1927). *Scientific Thought*. London: Routledge and Kegan Paul.
- Broadbent, Daniel E. (1958). *Perception and Communication*. New York: Pergamon Press.
- Campbell, John (2002). *Reference and Consciousness*. Oxford: Clarendon Press.
- Chalmers, David (1996). *The Conscious Mind*. New York: Oxford University Press.
- Chisholm, Roderick (1957). *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Clark, Austen (forthcoming). Phenomenal properties: Some models from psychology and philosophy. *Philosophical Issues* 18: "Interdisciplinary core philosophy", 2008.
- Cowey, A., and Stoerig, P. (1991). The neurobiology of blindsight. *Trends in the Neurosciences* 29: 65-80.
- Dehaene, S. & Nacchache, L. (2001) Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition* 79: 1-37.
- Dennett, Daniel (1988). Quining qualia. In A. J. Marcel & E. Bisiach, (eds) *Consciousness in Contemporary Science*. Oxford University Press, 1988, 42-77. (Reprinted in Block, Flanagan, & Güzeldere 1997, 619-642).
- Dretske, Fred (2006). Perception without awareness. In T. S. Gendler & J. Hawthorne (eds), *Perceptual Experience*. Oxford: Oxford University Press: 147-180.
- Driver, Jon & Vuilleumier, Patrik (2001). Perceptual awareness and its loss in unilateral neglect and extinction. *Cognition* 79: 39-88.
- Fodor, Jerry A. (1986). *Modularity of Mind*. Cambridge, MA: MIT Press.
- Freud, Sigmund (1900). *The Interpretation of Dreams*. Translated by A. Brill (1913). New York: MacMillan.
- LaBerge, David. (1995). *Attentional Processing*. Cambridge, MA: Harvard University Press.
- Nagel, Thomas (1974). What is it like to be a bat? *Philosophical Review* 83: 435-50. (Reprinted in Nagel 1979a, 165-80; and in Block, Flanagan, & Güzeldere 1997, 519-527)
- Nagel, Thomas (1979a). *Mortal Questions*. Cambridge: Cambridge

- University Press.
- Nagel, Thomas (1979b). Subjective and objective. In Nagel 1979a, 196-214.
- Neisser, Ulric (1967). *Cognitive Psychology*. New York: Appleton-Century Crofts.
- Pylyshyn, Zenon (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition* 80: 127-158.
- Raymond, J. E., Shapiro, K. L., and Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: an attentional blink? *Journal of Experimental Psychology: Human Perception and Performance* 18: 849-860.
- Raymond, J. E., Shapiro, K. L., and Arnell, K. M. (1995). Similarity determines the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance* 21: 653-662.
- Rey, Georges (1997). A question about consciousness. In Block, Flanagan, & Güzeldere 1997, 461-482. (Originally published in H. Otto and T. Uehling (1988) (eds.). *Perspectives on Mind*. Norwell: Kluwer Academic Publishers, 5-24.)
- Searle, John R. (1997). *The Mystery of Consciousness*. New York: New York Review of Books.
- Stoerig, P. and Cowey, A. (1989). Wavelength sensitivity in blindsight. *Nature* 342: 916-18.
- Stoerig, P. and Cowey, A. (1991). Increment threshold spectral sensitivity in blindsight: Evidence for color opponency. *Brain* 114: 1487-1512.
- Stoerig, P. and Cowey, A. (1992). Wavelength sensitivity in blindsight. *Brain* 115: 425-44.
- Weiskrantz, Lawrence (1986). *Blindsight: A Case Study and Implications*. Oxford: Oxford University Press.
- Weiskrantz, Lawrence (1997). *Consciousness Lost and Found*. Oxford: Oxford University Press.