PERCEPTION WITHOUT AWARENESS*

Fred Dretske

> We appear to be on the horns of a dilemma with respect to the criteria for consciousness. Phenomenological criteria are valid by definition but do not appear to be scientific by the usual yardsticks. Behavioral criteria are scientific by definition but are not necessarily valid.
>
> --Stephen Palmer in *Vision Science: Photons to Phenomenology (p. 629)*

Unknown to Sarah, her neighbor, a person she sees every day, is a spy. When she sees him, therefore, she sees him without awareness of either the fact that he is a spy or the fact that she sees a spy. Why, then, isn't the existence of perception without awareness (or, as some call it, *implicit* or *subliminal* perception) a familiar piece of commonsense rather thana contentious issue in psychology?[1] It isn't only spies. We see armadillos, galvanometers, cancerous growths, divorcees, and poison ivy without realizing we are seeing any such thing. Most (all?) things can be, and often are (when seen at a distance or bad light), seen without awareness of what is being seen and, therefore, without awareness that one is seeing something of that sort. Why, then, is there disagreement about unconscious perception? Isn't perception without awareness the rule rather than a disputed exception to the rule?

I deliberately misrepresent what perception without awareness is supposed to be in order to emphasize an important preliminary point--the difference between awareness of a stimulus (an object of some sort) and awareness of facts about it—including the fact that one is aware of it. When in the course of ordinary affairs S sees a spy without realizing he is a spy, S, though not

aware that he is a spy and, therefore, not aware that she sees a spy, is nonetheless aware of the spy. She sees him. She just doesn't know he is a spy. She can point at him and ask, "Who is that?" This is perception with awareness because S, though not aware of certain facts, is aware of the stimulus, her neighbor, the spy. Ignorance of the fact that one is seeing a spy does not impair one's vision of the spy, and it is lack of awareness of it, the spy, the stimulus, not the fact that he is a spy or the fact that one sees a spy, that constitutes perception without awareness.

Perception without awareness, unconscious perception, is therefore to be understood as perception of some object without awareness (consciousness[2]) of that object. Is this possible? Isn't perception defined as a kind of awareness? It is in some dictionaries. *Webster's Ninth New Collegiate Dictionary* tells us that perception is "awareness of the elements of the environment through physical sensation." If this is what it means to perceive the elements around one, then perception of these elements without awareness of them is like a bachelor getting divorced. It can't happen.

The scientific debate about unconscious perception is nourished by a variety of unusual--sometimes quite extraordinary--phenomena that are seldom encountered in ordinary affairs. We will look at some of these later, but, for now, it suffices to say that in deciding how to describe these results, scientists have found it necessary to distinguish different ways of perceiving x. If the way one person obtains visually mediated information about x is strikingly different from the way another person does it, then it proves useful to classify these as different ways of perceiving x. One distinction that has developed over the years is that between a conscious (supraliminal, explicit) and an unconscious (subliminal, implicit) perception of a stimulus: perception of x with awareness and perception of x without awareness of x. This sounds strange to ordinary ears--at least it does to mine--but the visual/cognitive deficits this language is used to mark are str ange

enough, or so it seems to many investigators, to justify these unusual descriptions. They have convinced many scientists that perception without awareness is a valid and useful concept. My purpose here is to look at these scientific studies for the purpose of learning what they reveal about consciousness. If psychologists can really identify something that deserves to be called perception without awareness, they must have an operational grasp on not only what it takes to perceive something, but on what it takes to be conscious of it. If this is really so, philosophers have something to learn from them.

1. Perception

If we are going to find a situation that deserves to be called perception of x without awareness of x we need to first agree about a test for perception (call it T_p) and a test for awareness (call it T_a) in which satisfaction of the first, though compatible with, does not require satisfaction of the second.

Many psychologists take transmission of information from x to S as an acceptable T_p. S saw x, despite not realizing it, despite denying she saw anything at all, if during the time she was exposed to x, S received information about x. If S's behavior is only (plausibly) explicable by assuming she was getting information about the F-ness (orientation, shape, color, location) of the stimulus, then although not aware that she perceived it, S must have gotten information about it. S therefore perceived it. Kanwisher's (2001: 90) description of perception as the extraction of perceptual information from a stimulus without assumption about whether or not this information is experienced consciously is typical.

Is this really a test? Is it a way of telling whether or not S perceived x? S can get (extract, receive) information about a stimulus and never use this information in any overt way. It might be

stored for later use and lost ("forgotten") before it can be used. Or maybe we haven't yet devised appropriately sensitive ways to show that S received information about x. If exposure to x makes a revealing difference in S's behavior, we can be sure S perceived x, but if there is no difference in S's behavior, we cannot conclude she didn't perceive x. Failure to manifest P is not a manifestation of not-P.

This problem can, perhaps, be minimized by remembering that we seek only a positive test, a behavioral sufficient condition, for perception. A behavioral necessary condition isn't really needed. If we can show that S got information about x —and sometimes this is perfectly obvious in S's behavior (if, for example, S describes x in glorious detail) --then even if it is sometimes difficult or impossible to show that S didn't receive information about x, we have what we need--a way of showing that S perceived x. All we need, in addition, to show that S perceived x without awareness is a test for awareness of x that S can fail while demonstrably getting information about x.

A more serious problem concerns the kind of information the receipt of which should count as perceptual. We don't want to include every reaction by S to x that carries information about x. People have quite distinctive allergic reactions to ragweed. Do these people, for that reason alone, perceive the ragweed? The rash on my leg carries information about my recent contact with or proximity to poison ivy. Should this physiological reaction to poison ivy count as perception of the plant? Human hair is a fairly sensitive indicator of relative humidity. Its length changes in regular and repeatable ways as the humidity varies. Do we perceive (with our hair, as it were) changes in humidity? If so, unconscious perception is a commonplace, not really worth special attention by psychologists and philosophers.

As these examples indicate, information about x may be deemed necessary, but if we trust commonsense judgments, it should not be taken as sufficient for perception. Something besides information about x is needed for perception of x. What might this be? Some will be quick to say *awareness* of x, a *conscious* experience of x, exactly the thing that is missing in the above examples. The reason allergic reactions to ragweed don't count as perception of ragweed is because they don't involve a conscious experience of ragweed. The physiological reaction carries information about the ragweed, yes, but it doesn't make one aware of the ragweed. So it isn't perception.

This move, at this stage of the game, would definitely settle matters: perception without awareness is impossible because awareness of x, a conscious experience of x, is required to make the receipt of information about x a perception of x. There are, however, other options. We might concede that perceptual experiences are necessary, but following Carruthers (2000: 147-79) allow for the possibility of unconscious experiences.[3] If we take this route, though, we should be ready to say what, besides carrying information about x, makes an internal state of S an experience of x? If a state needn't be conscious to be an experience, and allergic reactions to ragweed are not to count as experiences of ragweed, what additional properties of information-bearing states make them experiences of x? This is by no means an easy question to answer, but, guided by the way we conceive of conscious experiences of external objects, two additional requirements can be imposed: (1) the information in these states should be available for the control and guidance of action (if the experience is unconscious, of course, the actor need not be aware of this influence). (2) the information should be extracted from stimulation (as it is with conscious experiences) by accredited receptor systems.[4] If information about x is extracted from light by the photosensitive pigment of the retina, for instance, and this information is available for the control or modification

of behavior (reaching, pointing, grasping, identifying, describing), then the state (activity, event) carrying this information is a visual experience of x. Whether or not it is a conscious experience is a separate question. What these two additional conditions give us is the following: E is a visual (auditory, etc.) experience of x in S if E carries information about x, the information is extracted from light by photoreceptors in S's eyes (from sound by acoustic receptors in S's ears, etc.), and this information is directly[5] available for control of S's actions. Allergic reactions to ragweed don't count as perceptions (experiences) of ragweed because the information they carry fails to meet at least one--probably both--of these added constraints. The information isn't extracted by an accredited receptor system, and even if it is (according to some more liberal interpretation of "receptor system"), this information is not directly available for the control and guidance of S's behavior. So although S receives information about ragweed, she does not perceive it—not even unconsciously.

Even with the additional qualifications to come in a moment (see below), this "test" for perception of an object is not going to withstand philosophical scrutiny. Too many loose ends and philosophically troublesome qualifiers. Nonetheless, I propose to adopt it here. As I read the literature, this test comes reasonably close to the usage of people involved in this research while remaining tolerably close to ordinary language —close enough, perhaps, to justify using words like *see* and *hear* for what is being described. We could, I suppose, use subscripts to distinguish this special usage--if it is, indeed, a special usage--but as long as we keep clearly in mind what we are describing with these perceptual terms, confusion can, I hope, be avoided.

There is, however, another respect in which even if information arrives over an accredited sensory channel and influences a person's behavior, this is not enough for perception. I receive information about the Middle East, about continued violence in that part of the world, not by

seeing that part of the world, but by reading a newspaper in my living room. I don't have to see my gas tank to get information about it—that it is almost empty--through my eyes. What I actually see is the gas gauge, not the gas tank. Perception typically provides us with information about all sorts of objects we don't perceive. That is what instruments, radio, television, and newspapers (not to mention spies and informants) are for. They provide information about things we do not, perhaps cannot, ourselves perceive.

What is needed here, of course, is some principled distinction between direct and indirect perception. The rough idea is that if information about x is obtained by getting information about y where $y \neq x$ (e.g., a measuring instrument, a newspaper, a pilot light), then as long as y is not a proper part of x, perception of x is not direct. It is indirect. If the information you get about x is embedded in information you get in a direct way about y, then it is y you perceive. You may come to know about and react to x, but it is y you see, hear, or smell. The information I get about the Middle East from the newspaper is derived or indirect information—information delivered via information about the newspaper--that the headlines *say* there is continued violence in that part of the world. Information obtained about objects directly in front of us is presumably not indirect in this way. When I read the newspaper in normal conditions, I see the newspaper not simply because I get information about it (I could get that by looking at a photocopy), but because the information I obtain about it is not embedded in information I get about some more proximal object.[6]

I will not try to supply the required definition of *direct* perception. It would take us too far afield. I simply assume a satisfactory account is available. If my own account (Dretske 1981) is deemed unsatisfactory (Haugeland 1996 thinks it is), the reader is free to supply his or her own. If perception—conscious or unconscious—of physical objects (newspapers, gas gauges, people) is deemed possible, some such account must be presupposed.

Finally, a word about *how much* information one must receive (in a direct way) about x in order to perceive x. It needn't be much. Some information is necessary, but it needn't be enough to identify x. One can, after all, see a gadzit at a distance in poor light. It looks the same as a variety of non-gadzits look at this distance and in this light—like a small speck on the horizon. About the only information one gets in these conditions is information about its location. One can point at it. One can keep one's eye on it. One sees Venus in the night sky. It looks like a bright star. Without special instruments, information about the planet is not sufficient to identify it as a planet, not enough to distinguish it from a star. Nonetheless, one still gets information about it, information about its relative location in the night sky. In these circumstances, that is enough--enough, that is, to see it. In other circumstances (examining a bug under a microscope) one gets information about details of x without necessarily getting information about where x is (though one might know, on other grounds, where it is--in the lab, on the slide).

Pulling these ideas together, then, the proposed test for perception looks like this:

T_p: S perceives x = S has a perceptual experience (in our special inclusive sense) that provides (in a direct way) information about x.

From a scientific standpoint, one of the merits of T_p is that it does not require or presuppose consciousness. So it avoids the vexing issues we are trying to defer until §2. It leaves open the possibility of perceiving something unconsciously—without awareness of it. Whether that is really possible depends, of course, on whether an acceptable test for awareness can be formulated that makes awareness of x something more than perception of x and, therefore, something possibly absent when a subject perceives x.

2. Conscious Awareness

"Much of the long-standing controversial status of the study of unconscious processing revolves around the lack of a general consensus as to what constitutes an adequate operational definition of conscious awareness" (Reingold and Toth 1996: 159). An operational definition (at least an operationally useful necessary condition) of conscious awareness (of a stimulus) is our next topic. What can plausibly be used as a test for awareness of x that can fail when (according to T_p) x is perceived?

As already seen (spy example), we cannot use the fact that S does not believe[7] she sees an F (or believes she does not see an F), and therefore sincerely reports not seeing an F, to show that S is not consciously aware of an F. S might not know what Fs are. She might be confused or just not know about the existence (or prevalence) of Fs and, therefore, believe (and say) she is not aware of an one when she is, without realizing it, staring one in the face. Or the F S sees may be so far away, or exposed so briefly, or in such bad light, that it is impossible for S to tell (identify) what it is. That doesn't mean she isn't aware of an F. It only means she doesn't know what it is she is aware of.

More promising than identification or recognition is *detection* of a stimulus. To detect an F one doesn't have to know it is an F. One has only to be able to tell the difference (distinguish or discriminate) between it (the F) being there (wherever one is looking) and its not being there. One doesn't even have to be able to discriminate between Fs and non-Fs. Even if S doesn't believe in extraterrestrials and, as a result, refuses to believe she is seeing one (they look like large dandelions) in her front yard, S will be able to tell the difference between seeing one there ("Look at that huge dandelion!") and not seeing one there ("Its gone now"). If S can, in this way, tell the difference between the presence and absence of extraterrestrials (she needn't be able to

tell the difference between extraterrestrials and dandelions), then, no matter what she thinks they are, she can detect them. So, on a detection test of awareness, S is aware of extraterrestrials in her yard whether or not she knows (believes, judges, thinks) she is. She sees them consciously--with awareness. If, on the other hand, their presence makes no difference to S (they may, like electrons, be very tiny or move around too fast to be detected) then even when they completely surround her, she isn't aware of them. If she nevertheless gets information about them (directly), if she (in our extended sense) experiences them, then according to T_p she perceives them without awareness. Her perception of them is unconscious, subliminal, or implicit.

But how do we tell whether S can distinguish (discriminate, tell the difference) between x's presence and its absence? Do we leave this up to S? If S sincerely says she can't tell the difference between the presence and absence of x, if, according to S, things look the same to her whether x is there or not, does that settle the matter? Why? Why should we leave this up to S? Maybe she really can "tell" the difference (and we could show this if we found the right way to probe S), but she doesn't realize she can. Maybe her standards for seeing something are too demanding. Maybe she is biased in some way or isn't able or doesn't like to report things she sees in her left visual field.

We are now entering murky territory, territory in which there is a di vergence of opinion amongst psychologists about what appropriate criteria are. We can't say, simply, that S can discriminate x's presence from its absence--hence, detect x--if x's presence makes *a* difference to S because this is equivalent to equating detection with perception of x. Making *a* difference to S is just a way of describing S as getting information about the presence of x. We may, in the end, want to declare unconscious perception a theoretical impossibility, but this seems too quick.

Many scientists prefer a *subjective* criterion for detection, a test in which S's judgments and consequent reports about what she (consciously[8]) experiences or perceives (when they reflect a genuine power to detect the stimulus) define what S consciously experiences. Cheesman and Merikle (1984, 1986) clearly opt for this test in establishing subliminal perception: if a subject, asked to say whether anything is present, believes she is just guessing, then the stimulus is below S's subjective threshold of consciousness. S isn't conscious of the stimulus whether or not she is getting information about it. Of those stimuli S perceives, she is conscious of the ones she believes she is conscious of and not conscious of those she thinks she isn't. Restricting matters to stimuli S perceives[9] in location L (on the screen, in her left visual field, to the right of fixation point) we can express this subjective (superscript "s") test for awareness (subscript "a") as:

$^{s}T_{a}$:	**If S thinks (sincerely says) she is aware of x in L, she is aware of x; if she thinks (sincerely says) she isn't aware of anything in L, she isn't aware of x**

Two important clarifications:

1. $^{s}T_{a}$ is expressed not as a single condition necessary and sufficient for awareness, but as a dual sufficient condition: one sufficient for awareness, the other for lack of awareness. There is no proposed equivalence between awareness of a stimulus and thinking you are aware it. For good reason. We do not want to say of someone who perceives x that she is conscious of x if *and only if* she believes (sincerely reports) she perceives x. That would make a judgment that you perceive x into a necessary condition for awareness of x. It isn't. Human infants and a great many animals, I am assuming, are conscious of things around them.[10] They see, hear, and smell things in the same conscious way you and I do, but they do not (need not be able to) think or say that they are aware of them. They are conscious, yes, but, lacking conceptual sophistication, they do not think they are. Nor do they think they are not. They don't have thoughts on this topic. There is awareness,

but no *acknowledged* awareness. Acknowledgment, though, isn't necessary for awareness. That is a level of understanding you don't need to be aware of things.

Consider my dog, Fido. When Fido sees food in his bowl in what (I am assuming) is a fully conscious way, he may know there is food in his bowl. He may even be trained to "report" (by barking, say) to this effect, but unless the conceptual prowess of animals is vastly underestimated, Fido doesn't think (judge, know) that he sees, doesn't think he is conscious of, doesn't believe he experiences, food in his bowl. That is what *I* think. I believe this as a result of Fido's behavior (wagging his tail, eating from the bowl, barking), but that isn't what Fido thinks.[11] All Fido thinks, if he thinks anything at all, is that there is food in his bowl. That doesn't mean he's not aware of food in his bowl. Fido's impoverished intellectual life doesn't mean he is perceptually deficient. It doesn't mean he isn't aware of his food. He just doesn't know he is aware of it. So ${}^{s}T_a$ doesn't tell us anything about Fido. Fido satisfies neither the sufficient condition for awareness of food in his bowl nor the condition for the absence of awareness. If we want to know whether Fido is conscious of the food in his bowl (we already infer from his behavior that he perceives it), we would have to use some other test. The same is true for humans (e.g., one-year olds) before they understand what it means to perceive something (hence, are unable to judge and report that they are or are not aware of things). The inapplicability of ${}^{s}T_a$ to very young children and animals, however, does not mean that it is not a suitable test for those of us who know what it means to be aware of things. All it means is that we should understand the test to be of limited scope. It is designed for subjects who understand and can make judgments about not only the objects they are aware of, but their awareness of them.

2. For awareness of x, ${}^{s}T_a$ requires (as part of the sufficient condition for awareness) more than perception of x. It requires that the information (necessary for perception of x) be embodied,

specifically, in S's beliefs or judgments about whether she is aware, whether she perceives, x. S might perceive x and, according to $^{s}T_{a}$, *not* be aware of x because the information (necessary for perception) fails to produce a belief, a judgment, that she is aware of something. It is not enough that S believes (or says) that there is something (which turns out to be x) in L in order to be aware of x. According to $^{s}T_{a}$, S must actually judge (and be prepared to say) that she *perceives* it. We shall, in a moment, look at an *objective* test for awareness, $^{o}T_{a}$, that requires, for awareness of x, merely that S reliably judge or say whether x is present. This objective test for awareness of x is a test a subject can satisfy while believing she perceives absolutely nothing at all. $^{s}T_{a}$ is stronger. It requires that the subject not only reliably "say" a stimulus is present when it is present (Fido can do this much by barking), but think she is aware of it (Fido, I assume, can't do this).

Scientists have objected to $^{s}T_{a}$ on the grounds that it places on the individual subject the responsibility for establishing the criterion of awareness. (Eriksen 1959, 1960; 292; Underwood and Bright 1996: 4; Merikle 1984). It transfers responsibility for defining awareness (saying when a subject is aware of x) from the investigator to the subject of investigation.

> . . . factors unrelated to awareness, such as demand characteristics and preconceived biases, may lead subjects to adopt a conservative response criterion and report null perceptual awareness even under conditions in which conscious perceptual information is available. Response bias represents a threat not only to the validity of the subjective report measure of awareness, but also to its reliability. In particular, variability in response criteria makes it difficult to compare reports of null subjective confidence across-subjects, or within-subjects across conditions. (Reingold and Toth 1996: 162)

Such considerations led Eriksen (1959, 1960) to reject a subjective test as an adequate measure of awareness. He suggested, instead, that awareness be operationally defined in terms of performance on tasks that are independent of the subject's judgments about what she is or is not aware of. He urged the use of a forced-choice discrimination measure. Don't ask S whether she is aware of x or, if you do, don't take her word for it. If S doesn't think she can tell whether

x is present, if she professes being unable to see anything at all (hence, according to $^{S}T_{a}$, being aware of nothing at all), she is asked ("forced") to choose anyway: is there something there or not? If S thinks she can't see anything, she is urged to guess. It turns out in some conditions that subjects who say and think they are guessing are nonetheless able to "tell" (in a statistically significant way) whether x is there. Their "guesses" are more often right than wrong. In a sense, then, these subjects are detecting x when they do not believe they can. Even when they believe they can't. If we adopt this *objective* measure of awareness, $^{O}T_{a}$, subjects can be conscious of a stimulus while thinking and sincerely saying they are aware of absolutely nothing at all. They are no longer authorities on whether they are conscious of something.

This is called an objective test of awareness because the results are independent of what subjects *think* they are aware of. S may think (and say) she sees absolutely nothing in location L (thus being unaware of x in location L by $^{S}T_{a}$) and, yet, have her choices (she thinks of them as mere guesses) about whether something is present significantly affected by the presence of x. The objective test is obviously less demanding, than $^{S}T_{a}$. If we use the objective test, it turns out that S is aware of more than under the subjective test. If the objective test for awareness is used, perception without awareness still occurs (see Boornstein 1992: 193-94, for a summary), but it occurs less often and is harder to demonstrate experimentally.

Reingold and Merikle (1988, 1990, 1992) have argued that the validity of this objective test (as a test for *awareness*) depends on the plausibility of assuming that only stimuli S is aware of influence her discriminative responses.[12] If stimuli S is not aware of can affect S's decisions about whether x is present in an objective test for awareness, the objective test is not really a test for awareness. We don't want a test for S's awareness of x to allow things that S isn't aware of to affect the results that indicates awareness. Otherwise S needn't be aware of x to be counted, by

such a test, as being aware of x. We want the test to be *exclusive* or *pure*, to admit only factors that a subject, in some recognizable sense, is conscious of. This, though, is an exclusiveness, a purity, there is no assurance an objective test satisfies. Why infer that S is conscious of x in any recognizable pre-theoretic sense just because S's decisions about whether x is present are affected by x if S's decisions (choices, guesses) about whether x is present can be affected by things S isn't (in any recognizable pre-theoretic sense) aware of?

Reingold and Toth (1996: 163) think this a " . . . devastating problem for an objective test of awareness." It certainly seems to be a problem. It has pushed people otherwise sympathetic with objective methods back toward a subjective criterion of awareness, back towards something like ${}^{s}T_{a}$. Many—and, as I read the literature, *most*--scientists and philosophers feel that, even with its problems, ${}^{s}T_{a}$ comes closer than any objective measure in capturing something like our ordinary, commonsense, idea of what it takes to be consciously aware of something.[13] As a result of such widespread agreement, a subject's judgments (and, if they make them, sincere reports) on what she is (or is not) aware of remain the standard measure of consciousness in scientific studies of consciousness.

The use of the subjective test, ${}^{s}T_{a}$, is not only supported by raw conviction about what we mean to be talking about when we talk about awareness of a stimulus, it is also confirmed by a variety of theoretical approaches to the study of consciousness. I mention only two:

(1) Higher Order Thought (HOT) theories of consciousness (Armstrong 1968; Carruthers 1989, 2000; Dennett 1978; Lycan 1987, 1992; Rosenthal 1986, 1990, 1991) maintain that S's experience of x is conscious if and only if, at some higher level, S is aware that she is having the experience. Lower order mental states (e.g., S's perceptual experience of x) become conscious (making perception of x conscious) by becoming the object of a higher order thought. This

theoretical orientation makes it natural—indeed, almost unavoidable--to use ${}^{S}T_a$ as a behavioral test for consciousness. What *makes* an experience conscious is thinking you are having it; so if you think you are not having an experience of x, that you are not seeing x, you can't (unless you contradict yourself) be having a *conscious* experience of x.

(2) Global Access Theories of consciousness (Baars 1988) identify conscious processes with those whose informational content is accessible to a wide variety of output systems. If the only output (response, reaction) a given piece of information controls is direction of a person's gaze or size of the pupils in her eyes, then that information is not *globally* accessible--hence, not conscious. But if the information is also available to control a variety of bodily movements--e.g. where S points, where S looks, what S says--then this information is globally accessible and, therefore, conscious. If this is one's view about what makes information conscious, then ${}^{S}T_a$ is a natural test for consciousness. If S doesn't think (hence, is not prepared to say) she is aware of something in location L, the presence of something in location L is information that, even if S receives it (and thus perceives x), is, arguably at least, not globally (enough) accessible to qualify as conscious.

I will, in the next section, question ${}^{S}T_a$'s validity, but for now, and for these reasons, I tentatively accept it as the standard measure of awareness. If we use it, then, together with T_p, there are a variety of experimental results that indicate the existence--indeed, prevalence (in certain unusual cases)--of unconscious or implicit perception. Today, perhaps the best-known of these is blindsight (Pöppel, Held, and Frost 1973; Weiskrantz 1997), a condition in which patients with a partial "blindness" due to lesions in the visual cortex report seeing nothing in their blind field (thus, according to ${}^{S}T_a$, lacking awareness of the objects that are there) while obtaining information about them (as revealed by forced choice). Such accurate performance

(indicating the receipt of information) accompanied by lack of awareness has been identified in many categories of neuropsychological impairment (for instance, *numbsense*, a tactile analog of blindsight: Rossetti, Rode, & Boisson 2001). Unless there are reasons for rejecting either T_p or ${}^{S}T_a$, then, the existence of perception without awareness, unconscious, subliminal, or implicit perception, is an established scientific fact.

3. Validity of ${}^{S}T_a$

Is ${}^{S}T_a$ a reasonable test for awareness? For lack of awareness? It says, in effect, that belief that one is not aware of anything is to be treated as infallible. If you believe you are not aware of anything, you can't be wrong since believing this is sufficient for not being aware of anything. We have already seen that most such beliefs are fallible. Beliefs that I am not (visually) aware of a spy or an extraterrestrial are obviously fallible: I can see them and think I'm not. It isn't just specific kinds of things: spies, extraterrestrials, and poison ivy. I can believe I'm seeing absolutely nothing at all on the CRT monitor (I think and say that the screen is completely blank) and be mistaken. I confusedly think the figures I (consciously) see on the screen are figments of my own overactive imagination.[14] If "something" on a CRT means "something physical" (what else could it mean?), why can't I be aware of something there and believe I'm not?

What we are asking, remember, is not merely whether *perception* of something can occur without awareness of it, but whether *conscious* perception can occur alongside a belief that one is aware of nothing. If the reader feels that this is simply not possible, well and good. For them ${}^{S}T_a$ will be close to a definitional truth. It captures part of what they mean in describing someone as aware of a stimulus, as consciously perceiving something. But others--I confess to being one of

them--are not so sure. So it is at least worth looking at some of the unusual cases that challenge the validity of ${}^{S}T_a$.

A. Split Brains.

The corpus collosum is a large tract of fibers connecting the two cerebral hemispheres. When it is cut (to relieve epileptic seizures), the two hemispheres can no longer communicate with each other and patients reveal (under careful experimental tests) remarkable deficits relating to the perceptual, cognitive, and linguistic functions housed in the two hemispheres. Here is an example (taken from Palmer 1999: 631):

> A split-brain patient, N.G., was presented with a fixation point in the middle of a screen. Once she fixated it, a picture of a cup was briefly flashed to the right of the dot [information goes to left hemisphere where linguistic functions reside: FD]. She was asked what she saw, and she replied, "A cup." On the next trial, a spoon was flashed to the left of the dot [this information goes to the right hemisphere where linguistic functions are largely absent: FD]. Again she was asked what she saw, but this time she replied, "Nothing." She was then asked to reach under the screen with her left hand [behavior of left hand controlled by right hemisphere: FD] and pick out the object that had just been shown in her left visual field but without being able to see the objects. She reached under and felt each object, finally holding up the spoon. When she was asked what she was holding, she replied, "A pencil" [verbal behavior controlled by left hemisphere: FD].

N.G. was receiving information (in the right hemisphere) about the spoon--that it was a spoon. How else explain her ability (this is no lucky guess; she consistently gets it right) to pick it out (with her left hand). So N.G. must (T_p) see the spoon. But is she aware of it? Does she have the same kind of experience of it--i.e., a conscious experience--as occurs when she reports seeing a cup in her right visual field? Well, if she is aware of the spoon, and we take her word for what she believes, she certainly doesn't believe, and in this sense is not aware, that she is aware of it. If we take her verbal report as an honest and reliable expression of what she believes, then we are forced to conclude that N.G. believes she does not see a spoon on the left. She believes she

sees *nothing* in that part of her visual field (left of fixation point). So, according to ${}^{S}T_a$, she is not conscious of the spoon: a striking instance of perception without awareness.

But why take this as an instance of perception without awareness rather than a striking counterexample to ${}^{S}T_a$? Why not say that N.G. (or maybe the right hemisphere of N.G.'s brain) is conscious of the spoon, but because of a severed corpus collosum this information is not being transmitted to the speech centers on the left in charge of reporting whether, and if so, what is being perceived? Normal subjects detect and name objects in the left visual field because, after initially projecting to the right hemisphere, information is relayed via the corpus collosum to the left hemisphere where the stimulus can then be named and described. N.G.'s commisurotomy (severance of the corpus collosum) makes this communication about the existence of conscious experience (in the right hemisphere) impossible, but why should that be taken to mean that there is no conscious experience there? Wouldn't that be like concluding that nothing is happening in Foggyville simply because a storm that knocks out communication facilities prevents our getting news of the events occurring there? Why not conclude with Palmer (1999: 632) that the most likely explanation of N.G.'s behavior is that each hemisphere is conscious of the object projected to it, but only the left hemisphere is able to talk about it? The right hemisphere has conscious experiences. It just can't say (and think?) so.

If this is the way we interpret this extraordinary situation, then ${}^{S}T_a$ is not a valid test of awareness. N.G. is having (in the right hemisphere) a conscious experience of a spoon despite believing (if we take her report as an expression of what she believes) she is aware of nothing at all.

It may be objected, though, that in these special circumstances we are no longer entitled to accept what is reported as an accurate indication of what is believed. Once we split the person

into cerebral hemispheres with distinct streams of conscious experience, as we are now doing, we ought also split the person into distinct loci of judgment or belief: a left hemisphere capable of judging and talking about the experiences it is having (as evidenced by the person' s ability to report seeing a cup when the cup is in the right visual field) and a right hemisphere that, even if capable of making such judgments, is incapable of expressing them verbally. If we do this, then we do not have a violation of ${}^{S}T_a$. N.G.'s right hemisphere (in contrast to N.G.) does not judge (at least we have no evidence that it judges) that it is not aware of anything. N.G.'s right hemisphere may be like a human infant or an animal --aware of things but unable to report that it is (when it is) or that it isn't (when it isn't). We can take N.G.'s choice of a spoon with the left hand as a way of non-verbally reporting that there was, earlier, a spoon on the left, but this isn't yet a report that she was *aware* of a spoon (or anything else) on the left, and it is this report, a report (judgment) that she was aware of something on the left, that ${}^{S}T_a$ requires for awareness of the spoon on the left (see footnote 11). If we proceed in this way, ${}^{S}T_a$ renders no verdict on whether N.G.'s right hemisphere is conscious of anything just as it is silent on whether dogs and turtles are aware of things. We cannot say, at least not on the basis of T_p and ${}^{S}T_a$, whether this is perception with or without awareness. It is perception, yes: the right hemisphere is getting information (presumably in a direct way) about the spoon through the eyes and this information is controlling N.G.'s (left - handed) behavior. There is a visual experience (in our extended sense) in the right hemisphere, then, but is this experience conscious? ${}^{S}T_a$ does not tell us. It renders no decision about whether awareness is present in a location (the right hemisphere) where (as far as we can tell) no relevant judgments are occurring.

Whether split brains constitute a violation of ${}^{S}T_a$, then, depends on how one interprets the experiments. Is this to be understood as a person being conscious (with her right hemisphere) of a

stimulus while believing (according to her verbal report) she is not aware of anything? If so, we have a violation of ${}^{s}T_a$. Or is it a person's right cerebral hemisphere being conscious of a stimulus while (as far as we can tell) not believing that it is (or that it isn't)? If so, there is no violation of ${}^{s}T_a$. How do we choose between these two interpretations? I don't know. Maybe it is best to wait to see if anything else can be said against ${}^{s}T_a$.

B. Change Blindness

Change blindness refers to the inability of subjects to detect visible--sometimes quite prominent--differences. If the change (producing these differences) occurs when the subject can't see it (e.g., during an eye saccade) or during a suitable distraction (e.g., a mud splat), the differences produced by this change are sometimes hard to detect. Pictures of people standing around a jet airplane differ in a certain obvious (once you notice it) way: one picture has one of the jet engines (a prominent part of the picture) missing. Shown these pictures in alternating sequence (with a suitable intervening mask) subjects have trouble seeing the difference.

To have a familiar sort of example to focus on, suppose S looks at a scene in which there are seven people gathered around a table. Each person is clearly visible. S gazes at the scene for several seconds, runs her eye over (and, in the process, foveates [15]) each person at the table, but pays no particular attention to any of them. She then looks away. While S is looking away, an additional person--call him Sam—joins the group. Sam is clearly visible. There are now eight visible people. When S looks back, she doesn't notice any difference. Having no reason to suspect that a change has occurred, S thinks she is looking at the same group of people. When asked whether she sees a difference in the scene between the first and the second observation, S says, "No."

This is an example of change blindness—a clearly visible difference that S doesn't see. At least she doesn't believe she sees it. The question we need to ask is whether this is a genuine case of blindness. Did S actually see Sam on the second occasion (without noticing him) and, if so, was this conscious or unconscious perception? If S was consciously aware of Sam, she was clearly not aware (of the fact) that she was aware of him (or, indeed, of anything additional), but that, remember, is not the question. The question is not whether S is aware of *the fact* that she is aware of something different on the second occasion, but whether she is aware of the additional stimulus, Sam, on the second occasion. We are not interested in what *facts* S is aware of. We are asking what *objects* she is aware of.

It seems reasonable to say--at least it seems like a possible thing to say--that S not only saw Sam, but that her experience of him was of the same kind, a conscious experience, as was her experience of the other seven people. She was aware of Sam in the same way she was aware of each of the other seven people around the table. She was aware of the person who made a difference without being aware (realizing, noticing, or believing) that there was this difference. This is object awareness (of Sam, a member of the group) without fact awareness (that there is a difference in the group)--conscious perception of a stimulus without knowledge or realization that it is occurring. It would be completely arbitrary to say that S consciously sees only the same seven people she saw the first time and that, therefore, her perception of Sam, the new member of the group, is unconscious, subliminal, or implicit. Why just Sam? Why not each of the other seven people at the table? Or, despite S's protests (she thinks she saw, maybe, a half dozen people), isn't S aware of *any* people at the table?[16]

It seems more plausible to say that, on the second occasion, S perceived Sam in the same way S perceived every other person in the scene, the same way S perceived all seven people on

the first occasion--viz., consciously. S's experience on the second occasion was different from what it was on the first occasion, and the experience, on both occasions, was conscious. S just didn't notice (hence, remains ignorant of and, therefore, fails to report on) the difference. One can be conscious of the objects that constitute a visible difference and not be conscious of the fact that one is conscious of them.[17] If this is, indeed, the correct way to describe this situation, then ${}^{S}T_a$ is not a valid test of awareness. S is aware of an additional stimulus object--Sam--on the second occasion, something that makes a difference to what she consciously experiences, that she is not aware (does not realize) she is aware of. S believes there is no difference. S believes that there is *not* another person at the table and, therefore, that she is not aware of another figure. She believes her (conscious) experience, at least with respect to number of figures she sees, is the same. She is wrong. Contrary to ${}^{S}T_a$, S's beliefs and sincere reports about her own conscious experiences are mistaken.[18]

This conclusion doesn't mean that a person is always conscious of *all* the elements in a complex display. All it means is that one can be consciously aware of more than one realizes. If S looks (for several seconds) at seven hundred people in a room, seven thousand in a parade, or seventy thousand in a football stadium, she is unlikely to see everyone even if they are all clearly visible from her vantage point. The subjective impression of seeing hundreds, perhaps even thousands, of distinct elements, is (or may often be) an illusion.[19] These other (other than the ones attended to) objects may no more be seen (consciously or unconsciously) than are those objects whose image projects to the blind spot on the retina of a stationary eye (Dennett 1991). Maybe this is so, maybe it isn't. This is an empirical question. It depends, surely, on how much information about these objects one actually gets, and it is important to remember in this regard that not much information is needed to see an object. Information about color, for instance, is

clearly not necessary or we wouldn't be able to see things at dusk. Information about an object's shape isn't necessary. Things look pretty much the same at 600 yards, at acute angles, and when one isn't wearing one's glasses. That doesn't prevent their being seen. All that is really needed to see x is enough information about x to enable one to point at x and ask, "What is that?" If you get this much information about x, you see x, and, it seems reasonable to say (given that you can point at x and ask, "What is that?") that this is perception *with* awareness.[20] Our question about whether S is conscious of Sam, then, might be put this way: when S looks the second time, is her experience of the group such that she could have wondered "Who is *that*?" where "*that*" refers to Sam? Since this seems clearly possible with other members of the group, why not for Sam? If it is so for Sam, then her perception of Sam is conscious perception even when she insists she sees no difference.

If this is the correct way to think about change blindness--at least some instances of change blindness--then what change blindness shows is that sometimes--perhaps often--we do not notice some of the things we are consciously aware of. It shows that conscious experiences of the world are sometimes richer, more variegated, more textured, than the judgments one ends up making on their basis. It shows that "change blindness" might more correctly be described as a kind of "change amnesia" (Wolfe 1999: 74-75), an inability to retain (for judgment and report) information that is consciously registered at the perceptual level. It shows (Potter 1999: 35) that only some of the information consciously registered is still available when the time for action (e,g. reporting what is seen) arrives. It shows that ${}^{s}T_{a}$ is not a valid test for awareness of a stimulus.

C. Attentional Deficits: Unilateral Neglect and Extinction.

> Unilateral spatial neglect is a relatively common and disabling neurological disorder after unilateral brain damage. It is characterized by a lack of awareness for sensory events located towards the contralesional side of space (e.g., towards the left following a right lesion), together with a loss of the orienting behaviors, exploratory search and other actions that would normally be directed toward that side. Neglect patients often behave as if half of their world no longer exists. In daily life, they may be oblivious to objects and people on the neglected side of the room, may eat from only one side of their plate, read from only one end of a newspaper page, and make-up or shave only one side of their face. (Driver & Vuilleumier 2001: 40)

In *extinction*, unlike full-blown neglect, a patient can see and report on objects on either side of the visual field without difficulty if these objects are shown one at a time. If shown two objects at once, though, one on the good side and one on the bad side, they will report seeing only the one on the good side. Experiences of objects on the good side *extinguish* experiences of objects on the bad side.

If we use ${}^{s_r}T_a$ as our criterion for when a subject is conscious of a stimulus, then, since neglect patients insist they are not conscious of the objects they neglect (Driver and Vuilleumier 2001: 45) we must conclude with Palmer (1999: 637) that even if (according to T_p) neglected objects are perceived (there is evidence that they are[21]) they are not consciously perceived. The question we are asking now, though, is whether this is the right conclusion to draw. Could it be that these people are conscious of objects they say they don't see?

The reason one may be suspicious of the conclusion that patients are not aware of objects they neglect is that unlike blindsight where there is damage to the primary visual cortex, many of the neural pathways normally associated with conscious perception (including primary sensory areas) remain intact in neglect patients (Driver and Vuilleumier 2001: 45). Furthermore, unlike blindsight, some neglect patients can report an isolated light wherever it appears in the visual field

on either the good side or the bad side. What they can't do is report it on the left when there are (as there usually are in daily life) competing objects that "extinguish" it on the right. Nonetheless, despite the absence of report (and, presumably, belief) fMRI studies show that these extinguished objects on the left continue to activate the primary visual cortex and early extrastriate visual areas of the brain's right hemisphere in a manner similar to objects (when there is no competition from objects on the right) consciously seen on the left. As a result, Driver and Vuilleumier (2001: 66) are left wondering how these patients can possibly fail to be aware of the objects they neglect.[22]

There is, furthermore, substantial processing of information from "extinguished" objects beyond the primary visual cortex. Patients who report seeing nothing on the left can nonetheless make accurate judgments (they think they are guesses) about similarities and differences between objects presented simultaneously on the left and the right (Driver 1996: 200; Rafel and Robertson 1995). When asked to report where objects appear (on the left, right, or both sides) right-parietal patients (who extinguish on the left) report seeing nothing on the left, but when they are asked to count these same stimuli (are you seeing a total of one, two, or four objects?), some patients had no difficulty including the extinguished objects on the left in the reported total (Vuilleumier and Rafal:1999). We could, in accordance with T_p and ${}^{s}T_a$, describe these exploits as striking cases of perception (subjects are clearly getting information about extinguished stimuli) without awareness (i.e., while thinking and saying they don't see them), but the question now being asked is whether this is the *best* way to describe these results. Why not say, instead, that, contrary to ${}^{s}T_{a,}$ these subjects consciously experience objects on the left, but, like certain cases of change blindness, do not realize (notice, believe) they are experiencing them and, hence, do not report it. Neglect and extinction, after all, are typically classified as impairments of attention, and one can see things, consciously see things, one doesn't attend to.[23] That, at least, is what our discussion of change

blindness indicates. On this way of describing things, neglect patients are conscious of objects on the left. They just cannot attend to them in the way required for judgment and report. Their attention is directed only to objects on the right. Since objects in unattended regions cannot be reported (Dehaene and Naccache 2001: 8; Mack and Rock 1998), that would explain why subjects do not report seeing them. Besides an unshakable commitment to ${}^{S}T_a$, why take the further step of denying awareness of these objects?

This way of describing the results is not mandatory, of course, but we are not now looking for the way these patients *must* be described. We are asking how *best* to describe them. Since discussion of both split brains and change blindness has opened up the possibility--indeed (in the case of change blindness) the plausibility—that people are aware of things they believe they are not aware of, this seems like an available option in the case of neglect also.

4. An Alternative Test: Intentional Action.

If we forsake ${}^{S}T_a$, though, what takes its place? If we cannot rely on a subject's sincere reports about her own conscious experience, what can we appeal to? This isn't just a problem for the study of consciousness in others. It isn't just a problem of other minds. Giving up ${}^{S}T_a$ seems to create an epistemological problem about one's own conscious experience.

Blindsight may provide a clue. Blindsighters perceive things (in our T_p sense) without (unless told) realizing they perceive them. They normally believe they do not perceive objects in their scotoma (the "blind" portion of their visual field). They also exhibit striking cognitive and behavioral deficits with respect to these objects. If forced to guess or choose, they can more often than not correctly "say" whether x is vertical, blue, or moving (thereby exhibiting perception of

x), but they seem unable to exploit these facts in order to initiate spontaneous behavior. Saying (or believing) that x is vertical is not for them what it is for us, those of us who are conscious of x, something they are motivated or have reason to say. If they are cooperative, then, when asked to guess, they have reason to say something, but they do not, as we do, have reason to say vertical rather than horizontal. This suggests that perception of x with awareness might be distinguished from perception without awareness on the basis of the rational motivation for responses to x.

Something like this reasoning has led a number of investigators (e.g., Marcel 1983: 238-41; Dehaene and Naccache 2000: Flanagan 1991; 11; Hurley 1998: Chapter 4; Clark 2001; Milner and Goodale 1995; Van Gulick 1994) to view consciousness as, somehow, bound up with or manifested in rational agency. The idea, roughly, is that sensory information comes in two flavors: (1) in a form that makes it available for fixation of belief, rational planning, and choice. This is conscious experience. (2) In a form that although not available for planning, decision-making and grounding of judgment, can be used to control and tweak behaviors that have been rationally selected on other grounds. In the first form information helps determine what we choose to do and what we believe. It has a reason-giving role. In the second form information, though not available as a reason to do (or believe) anything, helps to determine how we do whatever we (with or without reasons) choose to do (Clark 2001; Jacob and Jeannerod 2003). This kind of picture also lies behind efforts by Dretske (1981, 1995), Evans (1982), and Tye (1995), to conceive of perceptual experience (each is talking about *conscious* perceptual experience) as that portion of incoming information available to cognitive centers for fixation of belief (reasons to believe) and goal selection (reasons to do).

To make this kind of picture plausible, especially if it is to be couched in terms of reasons, we have to carefully distinguish explanatory reasons, the reasons *why* S does A or believes P,

from justifying reasons, S's reasons *for* doing A or believing P, the reasons S (if able) might give to justify doing A or believing P. Explanatory reasons are those facts that explain, or help explain, why something happens—including why people (or animals) do (or believe) the things they do. As such, explanatory reasons have to be true. You can't explain why Sarah takes her umbrella by citing heavy rain if, in fact, it isn't raining. If it isn't raining, *that* can't be the explanation, the reason *why* she is taking her umbrella. Justifying reasons, on the other hand, needn't be true. Even if it isn't raining, that it is raining can be S's justification, the reason she has and the reason she gives, for taking her umbrella. If she (justifiably) thinks it is raining, or if it sounds to her like it is raining, that it is raining can be the reason she has—and certainly the reason she gives--for taking her umbrella whether or not it is raining.

Explanatory reasons have to be true, justifying reasons needn't be true. Justifying reasons needn't be true because a justifying reason (as I am using the term here[24]), unlike an explanatory reason, is the way the world, in either thought or experience, is represented to be. Since things do not have to be the way they are represented, since things are sometimes misrepresented, that it is raining can be S's reason for taking her umbrella even when it isn't raining. This "fact" (in scare quotes to indicate putative fact), qualifies as a reason for S to take her umbrella, even when it isn't raining, even when it isn't a fact, because from S's point of view it is represented as fact, as though it were true. It is, therefore, something whose falsity in no way diminishes an agent's rationality. From the agent's standpoint, a putative fact that is represented, in either thought or experience, to be a fact is, for purposes of both justification and motivation, as good as a real fact. A sunny day disqualifies the putative fact that it is raining as an explanation of anything, and therefore disqualifies it as the reason *why* S takes her umbrella, but it doesn't disqualify it as S's reason *for* taking her umbrella.

Although we cannot use what S believes--that it is raining--to explain why S believes P or does A when it isn't raining, we can use the fact that she believes it is raining or the fact that it sounds (looks, smells) like it is raining in our explanations of why S takes her umbrella. We can do this because unlike the putative fact *that it is raining*, that she believes it or that it sounds that way to her are genuine facts. She does believe it. It does sound that way to her. So although what she believes (that it is raining) cannot be the (explanatory) reason why she takes her umbrella, that she believes it is raining can.[25]

So much for the distinction between explanation and justification, the difference between reasons why S did A and S's reasons for doing A. To see how it plays out in an example more relevant to present concerns, and to anticipate its eventual use in a criterion for consciousness, consider the following situation. A subject in a psychological experiment is told to press the right button if the target, x, is vertical, the left button if it is horizontal. On the first trial S presses the right button. When asked why she pressed the right button, S (puzzled by the question since x is plainly visible to everyone) says, "Because x was vertical." Given her instructions, and being a cooperative subject, that was her reason, her justification, for pushing the right button. On the next set of trials exposure time is reduced. S is instructed as before: press the right button if x is vertical, the left if it is horizontal. S protests that she can no longer see x. She is asked to guess. She presses the right button. When asked why she pressed the right button, she says impatiently that it was just a guess. She had no reason. As she understands her instructions, she doesn't need a reason. She was, after all, instructed to guess and, for guessing purposes, one button is as good as the other. If we are convinced (by, say, a history of correct guesses at this exposure time) that S, despite her failure to realize it, perceives x (is receiving information about x in some primary way) we have, according to ${}^{s}T_{a}$, a case of perception (of x) without awareness (of x). Although

x's orientation is obviously influencing S's guesses (why else is she right so often?) its vertical orientation is clearly not S's justification or reason for pressing the right button. She doesn't have a reason for pressing the right (as opposed to the left) button. She is just guessing.[26] That is why she wouldn't venture a judgment un less "forced." That x is vertical, or that she perceives it to be vertical, may be the explanation of *why* she pushes the right button, but S herself has no reasons, no justifying reasons, for pushing the right button or, if she does, her reasons are not related to x's orientation.

Using our two tests, T_p and ${}^{s}T_a$, we have here a case of perception without awareness. We also have explanation in the absence of justification. There exists, in what S perceives, reasons why she pushes the right button without there being, in what she perceives, reasons for pushing the right button. This is suggestive. The obvious thought is that, perhaps, what constitutes a lack of awareness is not a *belief* that nothing is perceived—a failure, that is, of ${}^{s}T_a$, a test we have already found reasons to question—but the absence of justifying reasons. It is this idea that I would like to promote.

Aside from the differences between explanatory and justifying reasons already described, there is another striking difference: although F (some fact) can explain why S does A without S being aware of F, F cannot justify S in doing A unless S is aware of F. Justifying reasons, unlike explanatory reasons, are facts (or, in the case of justifying reasons, putative facts[27]) that one is necessarily conscious of. If S isn't aware that it is raining, the fact that it is raining cannot be her reason for saying or thinking that it is raining. It can't be her reason for taking her umbrella. That it is raining might still explain why (in a forced choice situation) S says it is raining. It might be the reason *why* she takes her umbrella. But it can't be S's reason *for* taking her umbrella—not if she is unaware of this fact.

If we combine this fact—the fact that one is necessarily aware of justifying reasons--with the plausible assumption that for those who perceive x, awareness of facts about x requires awareness of x itself, we get the following sufficient condition for awareness of a stimulus: S is consciously aware of x if for some action A or belief P the fact (or putative fact) that x is F is S's (justifying) reason for doing A or believing P. Conscious perception of x, perception of x with awareness of x, occurs when information (necessary for perception of x) becomes S's justification for belief or action. Call this principle J.

Notice, first, that although ${}^{S}T_{a}$ and J render the same verdict in many cases, they render it for quite different reasons. J is, or at least it is intended to be, less demanding than ${}^{S}T_{a}$. J does not require, as does ${}^{S}T_{a}$, that a subject believe she perceives x in order to be aware of x. All J requires is that information about x (that it is vertical, say) be S's justifying reason for doing (or believing) something, and this might be true of someone--a child or an animal, say--who lacks an understanding of what it means to be conscious and who cannot, therefore, satisfy ${}^{S}T_{a}$. As long as the animal or child can do things *for reasons*, as long as it can be motivated to act by having reasons to act, we can have grounds (according to J) for inferring that it is conscious of x even though it cannot think it is. If we are convinced that Fido, who just saw the cat run up this tree, has reasons for barking under this tree, and we are convinced, moreover, that his reason is that this is the tree (he saw) the cat ran up, then J tells us that Fido was conscious of the cat when it ran up this tree. As long as information about the cat--that it ran up this tree--is Fido's reason for barking here, Fido must not only have seen the cat run up this tree (and be acting on this information), he must have been aware of the cat when it ran up the tree. This is perception with awareness. J delivers this verdict for animals that lack an understanding of what it means to see things and, therefore, cannot (as required by ${}^{S}T_{a}$) judge themselves aware of a cat.

As this example is intended to illustrate, J does not require of the fact, F, that is S's reason for doing A that S believe (think, know, judge) that F. Belief, and the concepts it requires, is n't necessary. Maybe, as some philosophers think, Fido is incapable of having beliefs about cats and trees. If they are right, dogs don't think. They don't make judgments. If reasons for action, reasons for doing A, have to be states that deploy concepts, then, if these philosophers are right, dogs don't have reasons for the things they do. That doesn't mean they are not conscious, that they aren't aware of things. Fido might still see a cat run up a tree, but, lacking concepts, he couldn't see *that* (hence, know and believe that) a cat ran up the tree. If these philosophers are right (I don't think they are, but I set aside that quarrel for now), Fido is not barking under the tree because he thinks a cat ran up this tree. He is *caused* to bark there, yes. There are reasons *why* he barks there, yes. But he doesn't have reasons (a justification) *for* barking there. Not if justification requires belief. According to J, however, we can neatly avoid this deflationary (to animals) result without settling questions about the conceptual prowess of animals. Even if dogs lack concepts (and, therefore, beliefs) they can still have reasons for the things they do. Sense experience, if understood (at least in part) in representational terms, in terms that allow it to misrepresent surroundings, is enough. To have a reason for barking here, it is enough if Fido saw a cat run up the tree (one can do this without believing that a cat ran up the tree and without having the concepts CAT, TREE, and RAN UP). If Fido saw this happe n, and Fido's visual experience of this event represents it, in some phenomenal sense, as a cat running up this tree (or simply as something moving up something else), then that a cat (something) ran up (moved up) this tree (something else) can be the dog's reason for barking here. Fido needn't be able to conceptually represent what he saw as a cat running up a tree to have a reason, in what he saw--a cat running up a tree--for barking under the tree.[28]

If one has certain intuitions (see the discussion in §3) J also seems to do better than ${}^{S}T_{a}$ with split brains. Using ${}^{S}T_{a}$ we concluded that a commisurotomy patient who saw a spoon on the left was unconscious of it. Although N.G. picked out the spoon with her left hand, indicating thereby receipt of information about it, she denied seeing anything on the left. So, according to ${}^{S}T_{a}$, she was not aware of it. We found reason to question this verdict. Maybe there is, in the right hemisphere, a conscious experience of the spoon (there certainly is in the left hemisphere when she sees something on the right) without a judgment or report of this experience being possible because the hemispheres are separated. J justifies this result. It seems reasonable to say that the subject has a reason for picking out (with her left hand) the spoon. Selecting the spoon (as opposed to a knife or a cup) was not just a guess, a forced choice, or a random act. She carefully felt these objects before making her choice. If her choice is rational, something she has a reason to do, then although S may be incapable of verbally communicating her reasons to us, it seems right to say that S's justification for picking out the spoon was that it (the object she saw earlier) was represented—certainly in experience and maybe even in both experience and thought—as a spoon. That it was a spoon or that she saw a spoon is, therefore, her reason for choosing a spoon. If that is her reason, J tells us she was conscious of the spoon. Contrary to the verdict of ${}^{S}T_{a}$, this is perception of a stimulus with awareness of the stimulus by someone who believes (if we take what she says as an honest expression of what she believes) she is aware of nothing.

If we use J as a test (at least a sufficient condition) for a subject's awareness of a stimulus, we have a condition that, unlike ${}^{S}T_{a}$, not only yields acceptable results when applied to animals, young children, and split brain patients, but a test that explains why ${}^{S}T_{a}$, despite its drawbacks, is such an intuitively appealing test and why, in many cases, it works as well as it does. For if S is conscious

of x according to ${}^{S}T_a$, S will also be conscious of x according to the weaker, less demanding, condition J since the actions and perceptual judgments of cooperative subjects (the only ones we are considering here) will be reasonable, judgments and actions for which there will be (justifying) reasons. S's justification for thinking and saying she saw something in L will be--what else?--facts about the object she saw in L and this will include the fact that she saw it. So a subject conscious of x according to ${}^{S}T_a$ will also be conscious of x according to J. But not (as Fido and split brains indicate) *vice versa*. This suggests that maybe J is doing the heavy lifting. If S satisfies ${}^{S}T_a$, S must be conscious of x, yes, but not because she thinks she is conscious of x (Fido can be conscious of x without thinking he is) but because a fact about x (that she sees it) is the subject's reason for thinking she sees x, a reason she cannot have, according to J, without being conscious of x.

There are, however, problems with J that I have been ignoring. Aside from the fact that J is not really a genuine test (how, for instance, can one tell whether Fido has a *justifying* reason for barking under the tree) J expresses a sufficient, not a necessary, condition for awareness. We cannot use it, therefore, to establish what we set out to establish--conditions in which a person perceives x but is *not* aware of x. We can be certain that S is aware of x if S has, as her reason for doing A (or believing P), information about x, but if she doesn't do anything that has a fact about x as her reason, J is completely silent on the question of whether S is conscious of x. She might be or she might not be. S might be aware of the vertical stripes, but the fact that they are vertical may not be relevant to any of S's current plans or projects. There is nothing S is doing, or plans to do, or wants to do, for which their verticality is relevant. If we wanted to find out whether S was aware of these (vertical) stripes, we would have to arrange for S to do something for which their orientation was relevant. If S is cooperative, we could simply ask her whether x is vertical and assume that, if she gets it right (and isn't just guessing) the fact that x is vertical is her reason

for saying it is vertical. With animals and uncooperative subjects the task is harder. If S does nothing that has information about x as her justification, however, given only J, we are left to speculate about whether S is conscious of x. J is of no help. It certainly doesn't tell us that S is *not* aware of x.

To rectify this problem (it may be just papering over the problem) and secure a necessary condition for awareness we need to say that S is aware of x if and only if information about x is *available* to S as a reason. It is the availability of information for rationalizing and motivating intentional action (even if one is not capable of such action--e.g. paralyzed or buried in cement), not its actual use, that makes it conscious. We need something like the following (superscript "r" standing for a *reason* test):

$^{r}T_{a}$: S is aware of X = S perceives X, and information about X is available to S as reason (justification) for doing what she wants (chooses, decides) to do.

Securing a condition for awareness that is both sufficient *and* necessary comes at a price. How do we tell what information is *available* to S? To use an example from earlier, if S sees eight people around the table this time, but only seven people the first time, is information about Sam, the additional person, available to S if S denies seeing any difference? If S is really conscious of Sam without realizing it, how could we show that information about Sam is nonetheless available to her as a reason for doing or believing something when she denies seeing a difference? If this information is really available to her, S certainly isn't accessing it in answering questions about whether she sees a difference.

Demonstrating that this information is available might take experimental ingenuity, but, with the help of a little hindsight, one can imagine ways it might be done. We might, for instance, adapt Sperling's (1960) (also see Averbach and Coriell, 1961; Averbach and Sperling 1961) partial

report technique for determining what information is in a perceptual experience of a complex stimulus. In Sperling's experiments subjects are briefly shown (50 milliseconds) a set of letters:

T D A

S R N

F Z B

At this short exposure time, when asked to report as many letters as they can (the "whole report" condition), subjects identify at most four letters no matter how many letters are in the array. So if we took the number of letters they could identify as the number of letters they were aware of, we would have to conclude that, when exposed for 50 milliseconds to an array of nine letters, subjects were aware of, at most, four of them. Nonetheless, when asked[29], *after removal of the stimulus*, to identify the letters in only a single row (the "partial report" condition) subjects could often identify every letter in the row no matter which row they were asked about. In the partial report condition, then, subjects could identify *any* letter in the entire set despite being able to identify, at most, only four letters in the full report condition. This is not a case of attention being drawn to the queried row since the signal for which row to report occurs *after* removal of the stimulus. There is no longer anything out there (where the stimulus was) for their attention to be drawn to. Rather, subjects extract this information from what they describe as a conscious but rapidly fading image ("icon") that persists for a short time after removal of the stimulus. They use the information embodied in this conscious experience to identify letters in a stimulus that is no longer physically present. Sperling (1960: 20) concludes that: "A calculation of the information available to the Ss for their partial reports indicates that between two and three times more information is available for partial reports than for whole reports." If more information is available than subjects can use,

Sperling continues, they must choose a part to remember. In doing so, they choose a part to forget (23). Although this information is lost before being used, this information is nonetheless there, available to a subject (as revealed by the partial reports), in conscious experience, at the time (and shortly thereafter) the letters are seen. It is there, available as a reason to do (say) one thing rather than another. Sperling's brilliance consists of his finding the circumstances--partial rather than full report--in which information about each letter (though not information about all letters) could be used as a reason.

If we interpret these results as showing that subjects are perceptually conscious of more letters than they can (with such brief exposure) identify, that there is more information in their conscious experience of the letter-array than they can (in a "full report" mode) cognitively process and report on, we might use a similar procedure to demonstrate that a subject allegedly "blind" to differences is actually aware of the objects (e.g., Sam) that constitute the difference. After seeing the second (8 member) group, but before the "icon" (conscious experience) fades, a subject might be prompted by an arrow pointing at the position (formerly) occupied by Sam and asked, "Was anyone standing here?" If the answer is "Yes," it seems reasonable to conclude that the subject was, at the time she was looking at this group, consciously aware of Sam. If the same is true of the other seven people (this could be tested in the same way), we can conclude from this collection of "partial" reports that S was aware of all seven people on the first occasion, all eight on the second, but unaware (in full-report mode, as it were) that there was a difference in the n umber of people she was aware of. The availability of reasons and, thus (according to $^{r}T_{a}$), the difference in S's conscious experiences, is revealed by the "partial" reports, but not by the "full" or "whole" report. S cannot say, in full-report mode, whether there are differences in the scene, but by concentrating

on parts of the scene, we get the kind of difference in response that $^{r}T_{a}$ takes to be symptomatic of a conscious difference.

This, as I say, is only one possibility. There may be other ways of probing subjects to find out what they consciously experience. If we use $^{r}T_{a}$ as our guide, the way to go about determining what subjects are consciously aware of is not by asking them. Many of them (animals and infants) can't tell you. And of those who can tell you (adult human beings), many don't know. There is, often enough, too much going on in there (e.g., change blindness) for them to be very reliable about all that is going on in there. The way to proceed is, rather, by looking, in the most varied possible conditions, at what an agent finds it reasonable to do, at what, therefore, given suitable desires and circumstances, the agent is motivated to do. It is this information that most reliably indicates, and is the most accurate test of, how much of the world is being consciously experienced.

REFERENCES

Akins, K. 1996. Perception. Oxford University Press

Armstrong, D. 1968. *A Materialist Theory of the Mind*. London; Routledge

Averbach, E., & Coriell, A. S. 1961. Short-term memory in vision. *Bell System Technical Journal*, 40, 309-328.

Averbach, E. and Sperling, G. 1961. Short term storage of information in vision. In Colin Cherry, ed, *Information Theory*, pp. 196-211. London; Butterworth.

Baars, B. J. 1988. *A Cognitive Theory of Consciousness*. Cambridge University Press; Cambridge.

Beck, D.M., Rees, G., Frith, C. D., and Lavie, N. 2001. Neural correlates of change detection and change blindness. *Nature Neuroscience*, vol. 4, no. 6: 645-650.

Boornstein, R. F. 1992. Subliminal mere exposure affects. Pp. 191-210 in Boornstein & Pittman (1992).

Bornstein, R. F. and Pittman, T. S. 1992. *Perception without awareness: Cognitive, Clinical, and social perspectives*. Guilford Press, New York.

Block, N. 1995. On a confusion about a function of consciousness. *Behavioral and Brain Sciences, 18.2*. Reprinted in Block, Flanagan & Guzeldere (1997).

Block, N., Flanagan, O. and Güzeldere, G. eds. 1997. *Consciousness*. Cambridge, MA; MIT Press

Block, N. 2001. *Paradox and Cross Purpose in Recent Work on Consciousness*, pp. 197-219 in Dehaene, ed. (2001)

Brewer, B. 1999. *Perception and Reason*, Oxford University Press.

Campbell, J. 2002. *Reference and Consciousness*. Oxford University Press.

Carruthers, P. 1989. Brute experience. *Journal of Philosophy*, 86

Carruthers, P. 2000. *Phenomenal Consciousness*. Cambridge University Press.

Cheesman, J. and Merikle, P. 1984. Priming with and without awareness. *Perception and Psychophysics*, **36**, 387-95

Cheesman, J. and Merikle, P 1986. Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology*, **40**, 343-67.

Clark, A. 2001. Visual experience and motor action: are the bonds too tight? *Philosophical Review*, Vol 110, No. 4.

Coltheart, V., ed. 1999a. *Fleeting Memories*. Cambridge, MA; MIT Press

Coltheart, V. 1999b. Introduction: perceiving and remembering brief visual stimuli. Pp. 1-12 in Coltheart 1999a.

Coltheart, V. 1999c. Fleeting memories: summary and conclusions. Pp. 239-260 in Coltheart 1999a

Dancy, J. 2000. *Practical Reality*. Oxford University Press.

Dancy, J., Wallace, R.J., Darwall, S., et. al. 2003. Book Symposium on *Practical Reality* in *Philosophy and Phenomenological Research*, LXVII, 2, pp. 412-490.

Dehaene, S. ed. 2001. *The Cognitive Neuroscience of Consciousness*. Cambridge, MA; MIT press (A Bradford Book).

Dennett, D. 1978. Toward a cognitive theory of consciousness. In C. Savage, ed., *Minnesota Studies in the Philosophy of Science*, 9. Minneapolis, MN.; University of Minneosta Press.

Dennett, D. 1991. *Consciousness Explained.* Cambridge, MA; MIT Press.

Dienes, Z, and Perner, J. 1996. Implicit knowledge in peop0le and connectionist networks. Pp. 227-255 in Underwood, ed., 1996

Dixon, N.F. 1981 *Preconscious Processing*. New York; John Wiley and Sons.

Dorfman, J. Shames, V. A., and Kihlstrom, J. F. 1996. Intuition, incubation, and insight: implicit cognition in problem solving. In Underwood 1996: 257-296.

Dretske, F. 1969. *Seeing and Knowing*. University of Chicago Press.

Dretske, F. 1981. *Knowledge and the Flow of Information*. Cambridge, MA; MIT Press.

Dretske, F. 1993. Conscious Experience. Mind, Vol. 102.406, pp. 1-21

Dretske, F. forthcoming. Change blindness. *Philosophical Studies*

Driver, J. 1996. What can visual neglect and extinction reveal about the extent of "preattentive" processing? In Kramer, et. al. 1996.

Driver, J. and Vuilleumier, P. 2001a. Perceptual awareness and its loss in unilateral neglect and extinction. Pp. 39-88 in Dehaene, ed. 2001

Driver, J. and Vuilleumier, P. 2001b. Unconscious processing in neglect and extinction. *Out of Mind: Varieties of Unconscious Processes*, de Gelder, Haan, & Heykwood, eds. Oxford Press, 107-139.

Eriksen, C. W. 1959. Unconscious processes. In *Nebraska Symposium on Motivation, 1958*, edited by M. R. Jones, 169-227. University of Nebraska Press; Lincoln

Eriksen, C. W. 1960. Discrimination and learning without awareness: a methodological survey and evaluation. *Psychological Review*, **67**, 279-300.

Eriksen, C. W. 1990. Attentional search of the visual field. In D. Brogan, ed., *Visual Search*, pp. 3-19. New York: Taylor & Francis.

Flanagan, O. 1991. *The Science of the Mind*, 2nd ed. Cambridge, MA; MIT Press.

Gelder, B., Haan E., Heywood, C. 2001. *Out of Mind: Varieties of Unconscious Processes*. Oxford University Press.

Grice, P. 1961. The causal theory of perception. *Proceedings of the Aristotelian Society*, Supplementary volume 35: 121-152.

Haugeland, J. 1996. Direct Perception. In Akins 1996.

Holdender, D. 1986. Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: a survey and appraisal. *Behavior and Brain Sciences*, **9**, 1-23.

Humphrey, N. 1995. Blocking out the distinction between sensation and perception: Superblindsight and the case of Helen. *Behavioral and Brain Sciences, 18.2*: 257-58.

Hurley, S. L. 1994. Unity and objectivity. In *Objectivity, Simulation and the Unity of Consciousness, Proceedings of the British Academy*, 83 C. Peacocke, ed., Oxford University Press: 49-77

Hurley, S. L. 1998. *Consciousness in Action*. Cambridge, MA; Harvard University Press.

Jacoby, L. L. 1991. A process disassociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, **30**, 513-540.

Jacob, P. & Jeannerod, M. 2003. *Ways of Seeing: The Scope and Limits of Visual Cognition*. Oxford University Press

Kanwisher, N., Yin, C., and Wojciulik, E. 1999. Repetition blindness for pictures: evidence for the rapid computation of abstract visual descriptions. Pp. 119-150 in Coltheart, ed., 1999a.

Kanwisher, N. 2001. Neural events and perceptual awareness. Pp. 89-113 in Dehaene, ed., 2001.

Kramer, A. F., Cole, M. G. H., and Logan, G. D. eds. 1996. *Convergent Operations in the Study of Visual Selective Attention.* Washington, D. C.: APA Press.

Kunst-Wilson, W. R. and Zajonc, R. B. 1980. Affective discriminations of stimuli that cannot be recognized. *Science*, 207: 557-558.

Leibniz, G. W. 1981/1704. *New Essays on Human Understanding* (P. Remnant & J. Bennett, Trans & Eds.). Cambridge, England: Cambridge University Press

Lewis, D. 1980. Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy*, 58/3: 239-49.

Logothetis, N. K. 1998. Single units and conscious vision. *Proceedings of the Royal Society of London, Series B*, 353, 1801-1818.

Lycan, W. 1987. *Consciousness*. Cambridge, MA: MIT Press/A Bradford Book

Lycan, W. 1992. Uncertain materialism and Lockean introspection. *Behavioral and Brain Sciences*, 15.2: 216-17.

Mack, A. and Rock, I. 1998. *Inattentional Blindness*. Cambridge, MA; MIT Press

Marcel, A.J. 1980. Conscious and preconscious recognition of polysemous words: Locating the selective effects of prior verbal context. In R. S. Nickerson, ed.,*Attention and Performance VIII* (pp. 435-457), Hillsdale, NJ: Erlbaum

Marcel, A.J. 1983a. Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology*, **15**, 197-237

Marcel, A.J. 1983b. Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology*, **15**, 238-300.

Merikle, P. 1984. Toward a definition of awareness. *Bulletin of the Pssychonomic Society*, 22: 449-450

Merikle, P 1992 Perception without awareness: critical issues. *American Psychologist*, **47**, 792-5.

Merikle, P. and Reingold, E. M. 1992. Measuring unconscious perceptual processes. Pp. 55-80 in Boornstein and Pittman (1992).

Merikle, P., Smilek, D., and Eastwood, J.D. 2001. Perception without awareness. Pp. 115-134 in Dehaene, ed., 2001

Milner, D. and M. Goodale. 1995 *The Visual Brain in Action*. Oxford: Oxford University Press.

Moran, R. 2001. *Authority and Estsrangement: An Essay on Self-Knowledge*. Princeton: Princeton University Press.

O'Regan, J. K. and Noë, A. 2001. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, **24**, 939-1031.

Palmer, S. E. 1999. *Vision Science: Photons to Phenomenology*. Cambridge, MA; MIT Press/A Bradford Book.

Peacocke, C. 1992. *A Study of Concepts*. Cambridge, MA; MIT Press.

Perky, C. W. 1958. An experimental study of imagination. In *Readings in Perception*, D. C. Beardslee and M. Wertheimer, eds. Princeton, NJ, 1958.

Pöppel, E., Held, R. & Frost, D. 1973. Residual visual function after brain wounds involving the central visual pathways in man. *Nature*, 243 (405), 295-96.

Potter, M. C. 1999. Understanding sentences and scenes: the role of conceptual short term memory. Pp. 13-46 in Coltheart, ed., 1999a.

Rafel, R. and Robertson, L. 1995 The neurology of visual attention. In M. S. Gazzaniga, ed., *The Cognitive Neurosciences* (pp. 625-648). Cambridge, MA: MIT Press.

Reingold, E. M. and Merikle, P. M. 1988. Using direct and indirect measures to study perception without awareness. *Perception and Psychophysics*, **44**, 563-75

Reingold, E. M. and Merikle, P. M. 1990. On the inter-relatedness of theory and measurement in the study of unconscious processes. *Mind and Language*, **5**, 9-28.

Reingold, E. M. & Toth, J. P. 1996. Process dissociations versus task dissociations: a controversy in progress. In Underwood (1996), 159-202

Rosenthal, D. 1986. Two concepts of consciousness. *Philosophical Studies*, 94.3: 329-359.

Rosenthal, D. 1990. A theory of consciousness. *Report No. 40, Research Group on Mind and Brain*. ZiF, University of Bielefeld, Bielefeld, Germany.

Rosenthal, D. 1991. The independence of consciousness and sensory quality. In *Consciousness*, 1991, edited by Villanueva, E. Atascadero, CA: Ridgeview Publishing Co.

Rossetti, Y, Rode, G., and Boisson, D. 2001. Numbsense: a case study and implications. In Gelder, Haan, & Heywood 2001: 265-292

Shiffrin, R. M., Diller, D., and Cohen, A 1996. Processing visual information in unattended location. Pp. 225-246 in Kramer, et. al. 1996.

Shapiro, K.L, and Luck, S. J. 1999. The attentional blink: a front-end mechanism for fleeting memories. Pp. 95-118 in Coltheart, ed., 1999a

Sperling, G. 1960. The information available in brief visual presentations. *Psychological Monographs*, 74, 1-29.

Underwood, G., ed. 1996. *Implicit Cognition.* Oxford University Press.

Underwood, G. and Bright J. E. H. 1996. Cognition with and without awareness. Pp. 1-40 in Underwood, ed., 1996.

Van Gulick, R. 1994. Deficit studies and the function of phenomenal consciousness. In *Philosophical Psychology*, G. Graham and L. Stephens, eds. Cambridge, Mass.: MIT Press, pp. 25-50

Vuilleumier, P. and Rafal, R. 1999. Both means more than two: localizing and counting in patients with visuospatial neglect. *Nature Neuroscience*, 2, 783-784.

Weiskrantz, L. 1997. *Consciousness lost and found: a neuropsychological exploration*. New York: Oxford University Press.

Wolfe, J. M. 1999. Inattentional amnesia. Pp. 71-94 in Coltheart, ed., 1999a

Wolfe, J. M., Klempen, N., and Dahlen, K. 2000. Postattentive vision. *Journal of Experimental Psychology: Human Perception and Performance*, 26.2: 693-716.

ENDNOTES

*Thanks to Tamar Gendler, Güven Güzeldere, Peter Graham, John Hawthorne, Ram Neta, and Mark Phelan for helpful comments and criticisms. I also received useful feedback from members of the philosophy department at California State University at Northridge.

[1] See Holender (1986) and commentaries; also Dixon (1981), especially chapter 9, for an account of the sometimes acrimonious debate.

[2] I think my usage is fairly standard but, to forestall misunderstandings, I take *awareness* (of a stimulus) and *consciousness* (of a stimulus) to be synonymous. I sometimes speak of being *consciously aware* of an object, but, given my usage, this is a redundancy. I also use vision (seeing x) as my chief example of perception (and perceptual awareness), but what I say is intended to apply to all sense modalities.

[3] Causal theories of perception (e.g., Grice 1961, Lewis 1980) typically portray perception of objects in terms of a causal relation between the object perceived and some perceptual experience of the perceiver. If experiences are necessarily conscious, then, of course, causal theories are incompatible with unconscious perception. One can keep the causal condition while conceding the possibility of unconscious perception only by acknowledging the possibility of unconscious experiences.

[4] What is "accredited" is pretty much up for grabs, of course, but one thing it is meant to suggest is that it is a receptor system that, given normal continuation, sometimes at least, gives rise to *conscious* experiences.

[5] The "directly" is needed here because S's allergic reaction to x might, for instance, result in an itch that S feels or a skin rash that S sees. Feeling the itch or seeing the rash might lead to behavior—scratching or heading to the medicine chest—but in this case the information about x has been re-embodied (re-coded) in an acceptable experiential form. The information in the physiological reaction that gives rise to the itch or the rash is not *directly* available.

[6] In Dretske (1981: 155-168) I argue that it is a mistake to think that information about physical objects (the gas gauge, the newspaper) is always indirect—carried by information we receive about (even) more proximal objects (retinal images? sense-data?). The suggested restriction on the way the information about x must be received (i.e., directly) in order for x to be perceived does not, I argue, preclude perception of familiar objects and events.

[7] I include a subject's judgments (beliefs, knowledge, thoughts) as a legitimate part of a test for awareness because I assume, for convenience, sincerity (subjects believe what they say) and cooperation (subjects say what they believe relevant). In my way of proceeding, then, a judgment (that P) is equivalent to a report (that P), a more or less obvious piece of behavior. If readers find this objectionable, substitute "report" for "belief" and make the corresponding adjustment in the evidential status (less direct and, therefore, less reliable) of the test result.

.

[8] We have to add this qualification because, given our awareness-neutral test for perception (viz., T_p), S might believe she perceives x in this sense by being *told* she perceives x. Blindsighters, in fact, come to believe they perceive objects in this sense in this way. This, clearly, isn't enough to become consciously aware of the things perceived.

[9] Unless S perceives x (gets *some* information from x), questions about S's awareness of x don't arise. If S gets no information about x, S cannot be aware of x.

[10] Not everyone accepts this (see Carruthers 1989, 2000: Chapter 7), but I think it is widely enough accepted to justify assuming it here without further argument.

[11] In a series of experiments Logothetis (1998) trained monkeys to, as it were, "report" on their own experiences. Different stimuli were presented to each eye and the monkeys were trained to report on which stimulus they experienced. Under these conditions (binocular rivalry) human subjects report experiencing not a superposition of the two images, but an alternating sequence--first, say, vertical stripes (presented to the left eye), then horizontal stripes (presented to the right eye). A monkey's "reports" also alternate in this way. I think, however, it would be a mistake to treat these "reports" as reports about the monkey's experience. Unlike human subjects who understand the difference between x, the object of experience, and their experience of it, a monkey only reports on *what* it experiences--"vertical" vs. "horizontal" stripes--not *that* it is experiencing these different things. We assume that when a monkey reports the presence of a stimulus, he is aware of the stimulus in something like the way we are. We assume, that is, that the monkey would not be able to report "vertical" lines unless he was aware of them, but this (the fact that the monkey is aware of vertical stripes) is an inference we make about what it takes for an animal to report "vertical stripes." It should not be confused with what the animal is reporting--much less judging, believing, or knowing.

[12] This is called the *exclusiveness assumption* or (Jacoby 1991) the *process-purity* assumption.

[13] A small sample: Kanwisher 2001: 101, note 2, Cheesman and Merikle 1986: 344, 363; Dehaene and Naccache 2001: 6; Dienes and Perner 1996: 231; Merikle, Smilek, and Estwood 2001: 125; Underwood 1996: vi; Potter 1999: 41; Dixon 1981: 187.

[14] Cheves West Perky (1958) induced subjects to believe that the colored shapes they were seeing (carefully projected on the wall by Perky) were actually their own imaginative constructions. They saw things on L (the wall) that they didn't believe existed on L (they thought the "images" they were conscious of were in their own mind).

Skeptics tell me that Perky's experiments have not been successfully replicated. I assume here that failure to replicate an experiment is not a demonstration--experimental or otherwise--that the results of these experiments are invalid or otherwise unworthy of credence. It depends on why the experiments haven't been replicated. How many people tried? Some experiments are terribly hard to perform. If we can think that things in our own mind (e.g., hallucinatory images) are, in fact, real objects, why shouldn't it be possible to think that real objects are in our own mind?

[15] To foveate means to bring the image of each person onto the fovea (the sensitive part) of the retina.

[16] This is basically the argument I gave in (Dretske 1993) for the conclusion that, contrary to higher order theories of consciousness, one need not be aware of differences in one's experiences for these differences to be conscious differences. Also see Dretske (forthcoming) and the distinction between epistemic and non-epistemic perception in Dretske (1969). The difference between object awareness (of objects that make a difference) and fact awareness (that there is a difference) is, I think, a more perspicuous way of putting the distinction between *phenomenal* and *access* consciousness (Block's 1995). In other words, I agree with Nicholas Humphrey (1995) that the sensation/perception distinction is—or it *should be*--the primary distinction. See Block (2001) for a similar treatment of change blindness.

[17] One must be careful to distinguish: (1) awareness of the difference--which, in my idiolect anyway, implies awareness of the fact that there is a difference; and (2) awareness of an object—in this case Sam--that makes or constitutes the difference. S can be aware of Sam, the object that constitutes the difference, without being aware (in the factive sense) of the difference that Sam makes.

[18] Hurley (1998: 156-57) has an interesting discussion of the possibility of miscounting the points on a star-shaped after-image. If this is possible, and I don't see why it isn't, it makes the same point I am trying to make in the text--that it is possible for there to be differences in (the elements of) of conscious experience of which we are not (fact) aware.

[19] Contrary to what O'Regan and Noë (2001) suggest, the real question about change blindness is not whether one sees (or is aware of) *everything* in a complex scene (surely not!), but whether one can see (and be aware of) *more* of the elements, *more* of the detail, than one realizes. If there is anything (not *everything*, but *anything*) one can be aware of without realizing one is aware of it, then blindness to differences (failure to realize, notice, or believe that there is a difference) shows absolutely nothing, by itself, about what objects one is aware of. Change blindness is blindness to facts (a cognitive deficit), not necessarily (although it may sometimes be explained by) blindness to objects (a visual deficit). Although failure to see an object can explain why one doesn't notice it, there are other explanations. I am aware of nothing in the change blindness literature (other than an implicit acceptance of ${}^{s}T_{a}$--the criterion now in question) that shows that the cognitive failures they exhibit (not noticing differences that are plainly visible) are produced by a failure to consciously see the objects and/or features that make up these differences.

[20] Brewer (1999: 44-45) and Campbell (2002: 133) argue (and I agree with them) that demonstrative reference requires awareness of that to which reference is made.

[21] E.g., when shown otherwise identical pictures of a house one of which has flames coming from the left side of the house, neglect patients will deny seeing any difference in the pictures, but they will consistently "prefer" the picture of the house that isn't burning (Palmer 1999: 637)--thus indicating that they are getting information about that part of the house (the left side) that they neglect.

[22] It is known that information from unattended (therefore, unnoticed and unreported) stimuli is often processed to a very high degree (Shapiro and Luck 1999; Kanwisher, Yin, and Wojciulik 1999; Potter 1999). Beck and colleagues. (2001: 649) report finding face-specific activity in the brain when subjects are "blind" to ("unaware of" according to ${}^{s}T_{a}$) the change this activity signifies, indicating that, in some sense, the brain is registering changes that subjects report no experience of. Also see Shiffrin, et. al. (1996: 226) discussion of the "flanker effect.

[23] Significantly, Driver and Vuilleumier (2001: 116) report that many researchers have attempted to resolve these paradoxical results by proposing that the deficits exhibited in neglect and extinction may be more “attentional” than “sensory.”

[24] I should register here a small departure from an account of reasons with which I am sympathetic. Basically I agree with Dancy (2000; also see Dancy, Wallace, Darwall, et. al. 2003) and Moran (2001: 128) that S's reason (for taking her umbrella, for instance) is the fact that it is raining, not the fact that S sees (or believes) that it is raining. It is what S perceives or thinks (that it is raining), the *content* of her perceptual or cognitive state, not her perceiving or thinking that it is raining, that is her reason. Believing that it is raining may be necessary for the fact that it is raining to be her reason (it, so to speak, *enables* the fact that it is raining to be S’s reason), but believing that it is raining is not, thereby, her (justifying) reason.

What about false beliefs? What if S takes her umbrella because she mistakenly thinks it is raining? Since S would still (given what she thinks) *give* as her reason the “fact” that it is raining, this “fact” is still, in my sense, a justifying reason. As I read Dancy, he would deny that this is a reason. Under normal conditions (i.e., when use of an umbrella is to protect one from rain) one has no reason to take an umbrella when it isn’t raining no matter what one happens to think. One might think one has a reason, but one doesn’t. Whether or not Dancy is right about this (I think there is a sense of “reason” in which he is) I merely note that I am using the notion of a (justifying) reason more inclusively—as what is believed or experienced whether or not wh at is represented as true is true. In my sense, that it is raining can be S’s reason for taking her umbrella even when it isn’t raining. If she thinks it is raining, or it sounds like it is raining, and S take her umbrella, in part, because things seem this way to her, then that it is raining is (among) her (justifying) reason(s) for taking her umbrella.

[25] On a causal theory of knowledge, when S knows it is raining, her belief that it is raining is caused by the rain. In this special case, then, justifying reasons are (remote) causes of behavior. S’s reason for taking her umbrella--that it is raining--is the cause of her belief that it is raining, and the belief (causally) explains why she takes her umbrella.

.

[26] Hurley (1998: 148) puts the same point in terms of intentions: "When you guess on cue about a stimulus you are not conscious of, you guess intentionally. But information about the very stimulus in question does not feature in the content of your intentional guess in the same way it does when you intentionally report a stimulus you are conscious of, or when you act on it spontaneously. If information is conscious, you can report or act spontaneously on it: you can have the background intention to push the lever if, say, a light flashes, and you can then push it intentionally just because the light has flashed. More generally, if information is conscious then you can form an intention whose content is provided in part by that information and act on it just for the reason that information provides. Conscious information is available as an effective reason for acting. This is not the case when you can only guess on cue; the information in question does not activate your intention, or provide your reason for acting intentionally. You do not have intentional access to the information you can only respond to by guessing on cue."

[27] I hereafter concentrate on facts, the content of true beliefs and veridical experiences. I will, therefore, drop the distracting "or putative fact." It is to be understood, however, that S's justifying reasons need not be true. In the ideal situation, S's justifying reason (that it is raining, for instance) will be the content of a true belief, a belief that is part of the explanatory reason why S behaves the way she does (takes her umbrella or says that it is raining), but things aren't always optimal.

[28] For a broadly compatible view of experience, and the possibility of there being reasons in experience without the experiences requiring (on the part of the person having the experience) the concepts we use in expressing or describing their content, see Peacocke's (1992) notion of *scenario content* (roughly, the non-conceptual content of a perceptual experience). Seeing a square (in the right conditions) can give S a perceptual reason to believe it is a square even if S doesn't know what a square is and cannot, therefore, exploit that reason in coming to judge that it is a square.

[29] Subjects were "asked" this after removal of the stimulus by a distinctive tone (a different tone for each row) that signaled which row they were to report.