

FRED DRETSKE

CHANGE BLINDNESS

I'm looking for a friend in a crowd. I can't find him. He sees me and waves. I then see him. As it turns out, I was, several times, looking straight at him, but before he attracted my attention by waving, I didn't see him.

I can't find the cursor on my computer screen. In order to locate it, I move the mouse. This causes the cursor to move, and I spot it immediately.

It is not only movement – a change of position – that attracts attention. Changes in size, color, orientation, shape, and onset-offset are also quickly perceived. Of the dozen lights you see as you look down a dark street, it's the blinking one you notice first.

Though we are, in this way, sensitive to change,¹ the differences to which change gives rise, differences in position, orientation, color, shape and size, can be difficult to detect if the changes producing them are not themselves seen. If I am shown, in succession, pictures of a crowd of people, it is hard to see the difference between these pictures if they differ *only* in the arm position of one of the people in the crowd.² The differences in arm position can be as large and dramatic as those produced by a person waving, but if one doesn't actually see the arm move, it can be very hard to see these differences in arm position.

Or think about the blinking light. Can you imagine seeing the difference in the night scene of a city street – a difference of, say, only one light among hundreds – if you don't see it blinking on and off?

This difficulty in detecting, visible – sometimes quite conspicuous – differences when the differences are viewed successively is called *change blindness* – a phenomenon that has attracted a lot of scientific attention. It also deserves philosophical attention. It reveals, in a striking way, puzzling features that a philosophical



theory of perception should respect and clarify. This is what I hope to do in this paper.

1. CHANGES AND DIFFERENCES

I have already alluded to the difference between changes and differences. I want to spend another minute on this distinction since it will be important for my purposes. There can be differences (between x and y at time t_1 , for example) without change, of course, but even when the difference in question is a difference between x at t_1 and x at some later time t_2 , the word “change” can be used to refer to either an *event*, something that happens to x at (or during) this time, or to the *difference* in x that such a change produces. Differences are not events, not things that happen at (or during) a time, although the states (of x) that differ exist at a time. I come back to the question of just what a difference is later. For now, an example may help to make the distinction I’m after.

The minute hand on an analog clock moves continuously, but it moves too slowly to be seen. During the course of a full minute the hand occupies visibly different positions (6° of arc is easily discriminated by people of normal eyesight), but this difference, though visible, is not a difference one is likely to notice in the normal course of events. If two clocks differ only in the time they register – a difference of, say, only one minute – it will be hard to detect the difference between them unless one is looking for it. Even then, without helpful clues (minute marks are helpful clues) it may be hard to see the difference. The second hand, however, moves faster. You can actually see its movement. Unlike the change in position of the minute hand, one can see the second hand change position. With the second hand one sees an event – a 6° movement of the hand – and so the resulting difference between early and later positions is easily detected. It is, in fact, hard not to perceive this difference when one sees the change (event) that brings it about since seeing the hand’s movement is (among other things) seeing it occupy, in temporal succession, visibly different positions.³ One also sees the minute hand successively occupy visibly different positions, but, unlike the second hand, one doesn’t see it move.⁴ There is more to seeing change than seeing a succession of visibly different states.

What I mean by change as an event is what you experience when you see the second hand move to a visibly different position. You also see the minute hand occupy these same visibly different positions, but in the case of the slower moving minute hand you do not experience its movement. You do not see the transition between visibly different positions. The differences in position are still visible, but the change between them isn't. Slowness is to events what size is to objects. If objects are small enough, you can't see them. If events happen slowly enough, you can't see them. Even though it happens before your very eyes, you can't see grass grow (to be distinguished, again, from seeing *that* it has grown). If one sees the transition between different conditions the differences are easy to spot. If we don't see this event, but are merely exposed to the differences it brings about, the differences, though still visible, are hard to see.

In situations normally used to demonstrate "change blindness" subjects are exposed to differences (in pictures or scenes viewed sequentially) without being allowed to experience the change (or apparent change) that gives rise to these differences. The experience of change (the sort of thing we experience when we see a second hand change position) is, in fact, deliberately suppressed by eliminating the "motion signals" that normally give rise to a perception of movement. Change (the event) is made invisible by, for instance, timing its occurrence during that split second the eye is executing a saccade (a time when the eye is effectively blind). Only the differences existing at the two fixation points remain visible.⁵ "Change" blindness, therefore, is more like seeing a difference in the position of a minute hand, a difference created by a non-visible change, rather than a difference in the position of the second hand, a difference created by a visible change. For this reason, change blindness is more accurately described as difference blindness. It does not involve a failure to see change. The change, in fact, is carefully concealed.⁶ It is, instead, a failure to see the differences that change produces. That, in fact, is why change blindness is psychologically interesting. It is a difficulty in detecting something – a difference – that, unlike the change producing it, remains visible.

For this reason I will speak of difference blindness instead of change blindness. I think this is probably what most people are

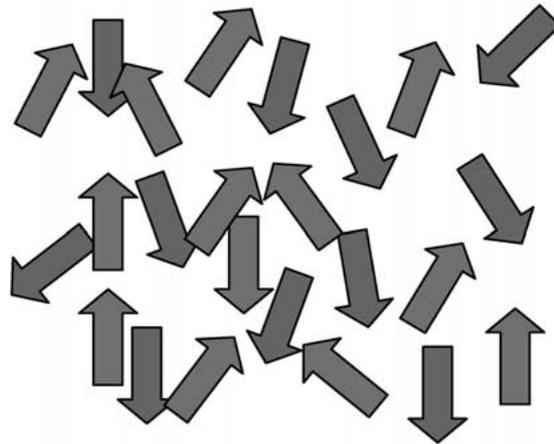


Figure 1.

talking about when they talk about “change blindness,” but I think it important to be explicit about this matter since the difference between seeing an event and seeing a difference is important. It underlies what William James (1890/1950, p. 747) was getting at when he said that only part of what we perceive comes through the senses; another part (James thought it was the larger part) comes out of our own head. Perceiving differences fits most naturally into that part of perception that comes out of our own head. I return to this point in the final section.

2. DIFFERENCE BLINDNESS

Compare Figure 1 and Figure 2. See the difference? Study the figures long enough (a few seconds should be enough) to assure yourself that you have seen every arrow in each array. If you are like me, several seconds will not be enough to see the difference. Though there is a visible difference (I now see it and my eyesight is no better than yours) you don’t see it. Failure to see something that is clearly visible to people of normal eyesight when you are looking right at it is – isn’t it? – a kind of blindness.

The orientation of a single arrow (call it Alpha) half way down on the far right is reversed. Alpha is pointing down in Figure 1, up in Figure 2. Once you are told about it (or notice it for yourself), it is easy enough to see the difference. You now know where – and, perhaps, how – to look.

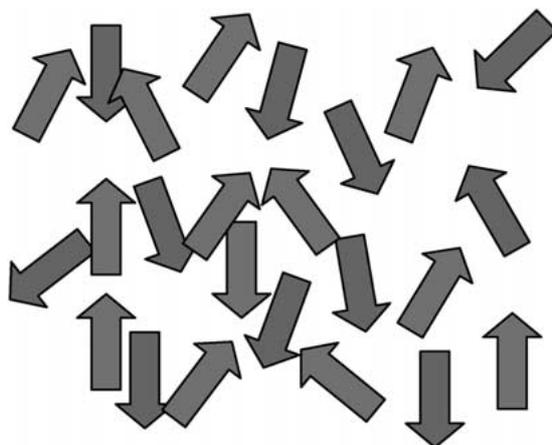


Figure 2.

You may feel that the difference between Figures 1 and 2 is so minor, so insignificant, that one could hardly be expected to see it. Perhaps it is. Nonetheless, we have trouble even when the differences are much more dramatic. Subjects have difficulty seeing the difference between pictures of two people when their faces are switched. They fail to see a difference between pictures when a prominently displayed parrot (it occupies over 25% of the picture) changes color. As long as the changes (= events) producing these differences are not visible, the resulting differences are hard to spot.

Grimes (1996, p. 92) describes a surprising instance of difference blindness that occurred in a study by McConkie and Zola (1979). Experimental subjects were asked to read text when the letters were printed in alternating case:

- (1) ThE sPaCe ShUtTIE tHuNdErEd InTo ThE sKy On A
cOIUmN oF sMoKe

Though difficult to read at first, subjects quickly became accustomed to it with practice. The text was displayed on a computer screen. Eye movements were monitored and during every saccade (an eye movement from one fixation point to another – something that happens two or three times a second) the case of every letter was changed. That is, (1) was replaced by (2):

- (2) tHe SpAcE sHuTtLe ThUnDeReD iNtO tHe SkY oN a
CoLuMn Of SmOkE

Then, during the next saccade, (2) was replaced by (1). The alternation continued so that the appearance of every letter was dramatically different on each new fixation. Since the change took place during a saccade when the eye transmits no information, subjects did not see the change (event) that occurred several times every second.⁷ That was no surprise. What was surprising is that subjects failed to notice a difference in what they saw on each fixation – the difference between (1) and (2). There was no awareness that anything unusual was going on. Although every letter they saw changed its appearance several times each second, they never noticed it.

In describing difference blindness I have been deliberately careless in my choice of words. I have spoken, indifferently, of spotting, detecting, seeing, noticing and perceiving differences. Are all these the same? Can you see changes you don't detect or notice? Can you see things you don't perceive? If you can, why describe a failure to detect (perceive, notice) a difference as a form of blindness? Isn't it just a failure to notice (detect, perceive) something you see. This isn't blindness, just inattention. But are differences the sort of thing one can see and not notice? And what, exactly, is it that one is supposed to be blind to when one fails to detect a difference? Is it the objects that differ (Alpha in Figures 1 and 2)? Or is it the properties of Alpha (its orientation) which makes Alpha (and, therefore, Figures 1 and 2) different? Or is it something else?

In order to avoid getting entangled in all these questions (I do get entangled in some) I will stick to the verb "to see." I avoid other, quasi-perceptual, terms. I especially avoid the word "detect" although this is a term that is often used in the scientific literature on change blindness. This word (and several others like it) darken the philosophical waters I am trying to navigate because it is seldom made clear how detecting something is related to seeing it and, therefore, why failure to detect something is a form of blindness. Since I am interested in whether we are blind to certain visible differences, and blindness is, I take it, an inability to see visible things, it seems proper to start the investigation by looking at what we see and what we don't see. We all know that *seeing* isn't trouble-free – quite the contrary – but at least we share some common intuitions about when this verb applies and why, and I hope to lean on these intuitions in order to get clear about what it means to see

differences and, therefore, on what it means to be blind to them. I will also concentrate on vision. Although I think the phenomenon of interest exists in every sense modality, I focus on difference *blindness*, not, say, difference *deafness* or difference *numbness*.

3. SEEING AND SAYING

Before jumping to conclusions about what differences people do and do not see, let's start with what people *say* they see and don't see. We can then worry later about what can be inferred from this.

In certain experimental conditions, people of perfectly normal vision say they don't see differences that other people (or themselves at a later time) find conspicuous. What can be inferred from this fact? Well, if they are sincere – and I will assume that sincerity is never a problem (everybody always believes what they say) – we can infer that these people do not believe they see a difference. We can infer more. We can infer not only that they lack a certain belief (that they see a difference), but that they have the belief that they don't see a difference.⁸ What, then, can be inferred from the fact that a person believes she does not see a difference? That she doesn't see a difference? For most values of *x*, *S*'s believing she does not see *x* is compatible with *S* seeing *x*. It is compatible with fully conscious perception of *x*. One needn't believe in flying saucers, miracles, and spies to see them. If they exist, you can see them while believing you are not seeing them. Skepticism about *x* makes one ignorant of, but not blind to, the *x*'s one sees.

Philosophers and psychologists who embrace “reportability” criteria of consciousness (and, hence, of conscious or explicit perception) are clearly not thinking of spies and flying saucers when they conclude that a subject is not aware of an *x* from her denial that she is aware of an *x*. One can perceive – *consciously* perceive – spies and flying saucers (teapots, bicycles, etc.) while sincerely denying awareness of any such thing. Behavioral measures of consciousness that tie a person's perception (and by “perception” I always mean conscious or explicit perception) of *x* in location *L* too closely to the person's ability to report his awareness of *x* in *L* tend to confuse conscious perception of objects (things we can perceive without noticing or recognizing in any way) with conscious perception of

facts – either the fact that there is an x in L or the fact that one is aware of an x in L . Awareness of facts requires some cognitive uptake, and therefore an ability (for those who possess language) to report awareness of an x in location L , but it is not at all clear that awareness of objects in location L requires a similar ability to report. Although you can't see (the fact) that there are spies in the neighborhood without believing that there are spies in the neighborhood, you can certainly see spies in the neighborhood while believing that there are none (and, therefore, that you are aware of none).

Talk of spies and flying saucers – the sort of thing one can see while believing one is not seeing it – may be brushed aside as irrelevant. Those sympathetic with “reportability” criteria of perception have in mind something quite different when they take, as authoritative, the word of a person denying perception of x . They have in mind something closer to what is called a “stimulus detection” measure of perception. The idea here is that a person does not see an x in location L if he or she sincerely reports (hence, believes) seeing absolutely nothing in location L .⁹ A person who believes there are no spies may not report seeing spies in the neighborhood, but she will report seeing something in the neighborhood – her neighbors, for example. And if the neighbors she reports seeing are, unknown to her, spies, then, contrary to what she believes, she sees spies. But if a person sincerely reports seeing nothing at all in the neighborhood, that should settle things. She doesn't see anything in the neighborhood and, in particular, no spies. No differences either.

But subjects in these experiments do not report seeing nothing at all. As the examples (Figures 1 and 2) illustrate, subjects are typically shown, and they will report seeing, quite a lot. What they report not seeing is a difference. So we come back to the question: is seeing a difference like seeing a spy (flying saucer, etc.) – something you can see without being aware you are seeing it – while believing, in fact, that you are not seeing it? Or is seeing a difference different? Are visible differences, unlike spies and flying saucers, the sort of thing that, when you see them, they jump out at you, grab you by the cortex, and proclaim their identity: *I am a difference*? Are they, then, unlike spies and miracles, the sort of thing you can't see without recognizing? Only if this is so should we conclude that people are blind to differences they sincerely profess not seeing.

4. SEEING DIFFERENCES

A friend recently shaved his moustache off. I didn't notice. I thought he looked different somehow, but, I had the vague impression that he had gained weight. His face looked fuller, rounder, younger, but I didn't say anything. Gaining weight is not something people like to have their attention called to. His face was visibly different from what it was the last time I saw him. Moustache then; no moustache now. Despite not noticing the (real¹⁰) difference – the absence of a moustache – did I, nonetheless, see it?

To be more specific (we'll get back to "the difference" in a moment), did I see the skin on his upper lip, the skin that was formerly concealed by his moustache, skin that I therefore did not see when he had a moustache? I saw him at close range in good light. We were sitting across from each other at a small table. If I didn't see the skin on his upper lip, what did this part of his face look like? What did I see between his nose and mouth? Nothing? Did his nose, then, appear to be directly above his mouth with nothing separating them? Wouldn't I have noticed that? If I didn't see that part of his face – his upper lip – why is that the *only* part of his face I didn't see? Or didn't I see *any* part of his face?

Given that my friend has a normal upper lip, it seems most plausible to say that his upper lip looked to me pretty much the way normal moustacheless upper lips look. If this is indeed so, then I saw something the last time I didn't see earlier – viz., my friend's upper lip. Since I didn't see it earlier (the moustache was concealing it) this means I saw something earlier – his moustache – that I didn't see the last time, and I saw something this last time (his upper lip) I did not see earlier. But the difference between a moustached upper lip and a moustacheless upper lip is (isn't it?) exactly the difference I am asking whether I saw. So I must have seen the difference. I just didn't notice it.

If we use this model – I'll call it the *object* model – to understand change blindness, we are driven to the conclusion that subjects (of normal eyesight, looking in the right direction, under suitable conditions of illumination, etc.) who say they see no difference are probably¹¹ wrong. They do see it. They just don't notice it. They just don't identify it – at least not as a difference. But they do see a difference. According to the object model, differences are like objects –

spies, for instance – things you can see without ever knowing you are seeing them. They may be hard to identify, perhaps, but not hard to see. Like spies, they are things you can be conscious of without being aware (conscious) that you are conscious of (perceive) them.

That is one way of describing what I saw. It is inspired by the fact that I'm not blind. When I saw my friend on the two occasions, I not only saw the visibly different conditions, I saw (on the one occasion) the thing (moustache) that made them different. What more could possibly be relevant to seeing a difference? The fact that I didn't notice the difference is what explains why I didn't comment on it and why, later, I might, reasonably enough, deny having seen a difference. But, just like spies, differences one reasonably denies having seen are not always, not necessarily, differences one doesn't see. They are sometimes differences we don't recognize.

This way of interpreting difference blindness is (or seems to be) supported by a vast array of familiar experiences. You are looking for your husband in a crowded marketplace. You can't find him. Later, after you've found him, and he tells you he was standing directly in front of the fruit stand you looked at several times, you, in effect, plead blindness: "I didn't see you." Wrong! You (probably) did see him. You just didn't recognize him. He was, after all, standing in your line of sight only a few yards away in sunlight. You didn't see through him. You surely didn't see the apples on the stand directly behind him. The reason you didn't see the apples is because he was in front of them. If he really did prevent you from seeing the apples, you must have seen him. How else can x block your vision of y ? The reason you didn't recognize him is that you were looking for someone in a brown sweater so you ignored the person you saw standing there in a blue sweater. According to an object model of seeing differences, differences are no different.

When we talk about concrete objects – moustaches, lips, and people – all this is obvious. At least it should be obvious. We can see such things without noticing or recognizing them. I do not have to notice, recognize, or identify every book I see on the shelves as I scan them in search of a particular book. I may, in this process, see hundreds, even thousands, of books. I notice, I pay attention to, only one or two – sometimes, when I'm being very focused, only the one with the distinctive binding I am looking for. If we

use this model to understand change blindness, it means we could – and probably do – see differences every moment of our waking life that we simply don't notice. If you examined Figures 1 and 2 attentively, then, according to the object model, you probably¹² saw not only Alpha, but Alpha's pointy side on the top in Figure 1 and on the bottom in Figure 2. You, therefore, saw the difference between Figures 1 and 2, and you saw it before you noticed (or were told) that there was this difference.

There is, though, another way of understanding change blindness. On this alternative model – I'll call it the *fact* model – differences are not another kind of object. They are not like moustaches, books, people, lips, and arrows – objects that can be seen without being noticed. I can see a moustache, a book, or a person, concrete physical objects, without noticing them, and these objects can change state in visible ways. The fact that I can see, without noticing, both *x* and the (different) properties of *x* (its redness and its greenness), though, does not mean that I can see, without noticing it, the *difference* between these states. For in referring to something as “a difference” we refer not to the objects or properties that make up the difference (e.g., the different colors of *x*), but to the fact that they make a difference, and facts (e.g., that *x* has changed color), unlike concrete objects (*x*) and properties (colors), cannot be seen without being noticed. The facts we perceive are just what we notice, what we come to know, about the objects we perceive. So if you perceive a difference you have to notice it. Seeing a fact *is* (among other things) noticing it. Unlike the moustache itself, if you don't notice the fact that he has a moustache, you don't see it. Difference blindness is blindness to facts, not objects. On the fact (but not the object) model you can be absolutely blind to differences and still see all the things (object and their properties) that make up the difference.

According to the fact model, difference blindness is not like color blindness. If I am color blind I don't actually see the colors that make a difference. The fact that I can't see the colors explains why I can't see a difference (in color) between red lights and green lights. If change blindness is, however, a failure to see a fact, then blindness to a difference in color is compatible with seeing the colors that make the difference. In change blindness, unlike color, blindness, it

is the difference in color, not the colors that differ, that one is blind to.

Difference blindness is like blindness to problems, solutions, answers, and a great many other things designated by abstract nouns. The abstract noun “stands in” for a fact. If the problem is the plugged drain, it doesn’t follow that you see the problem just because you see the object – the plugged drain – that is the problem. You have to see (the fact) that it is a problem. You don’t see the answer if all you see is a number (written on the board, perhaps) that is the answer. You have to see that it is the answer. My eyesight can be very good, much better than yours, yet you see the answer (problem, solution, etc.) and I don’t. In order to see what you see, I don’t need better glasses, I need better understanding.

If there are, then, these two ways of understanding the perception of difference – as the perception of objects that differ vs. the perception of the fact that they differ – which is the correct way to understand difference blindness? Are differences objects or facts?

The fact interpretation, I believe, has the better of it. When speaking of what can be seen, differences are like problems and answers – things that can’t be perceived without being known and, therefore, attended to. Just as I don’t see the problem if I am totally unaware (of the fact) that there is a problem, I don’t see a difference if I’m totally unaware (of the fact) that there is a difference. If we both see twenty-three books the first time and twenty-four the second time, and you notice that there are more books the second time and I don’t, you see something – a difference in the number of books – I don’t see even though, on both occasions, I see all (and only) the books you see.

Facts don’t reflect (or emit) light, objects do. Although not everything we see that depends on our deployment of our eyes reflects (or emits) light,¹³ they all, generally speaking, involve objects that do and perception of them depends on our perception of these associated objects. Facts aren’t like that. In a sense, facts aren’t visible. That doesn’t mean that they are invisible. Like numbers, problems, and answers, facts don’t play in that league.

One can’t see facts and not notice them, but that isn’t because facts are particularly conspicuous or salient objects, things that force their way upon our attention. It is, rather because the facts we

perceive, unlike the objects and events we perceive, are the “objects” (content) of propositional attitudes, intentional states like belief, judgment, and knowledge, that are constituted by awareness of the fact (or, in the case of belief and judgment, putative fact) on which the mind is directed. That is why seeing a difference you don’t notice is like thinking that P without noticing or attending to P. You can, of course, think something – e.g., that you are in Paris – without really noticing or being aware that you think this, but to think you are in Paris you have to be aware of what you think – that you are in Paris. That is part of what thinking is.

5. FACT PERCEPTION

The fact model of difference perception (and blindness) encourages the idea that difference perception is not a visual phenomenon at all. Echoing James, fact perception is that part of perception that occurs entirely “in the head.” Perception, real perception, the kind of perception that involves essential use of the eyes, the perception of objects and properties, is all over before fact perception even begins. On this way of thinking, fact (and, therefore, difference) perception is simply the conclusion-drawing phase of a perceptual process the premises of which are supplied by object and property perception. We experience objects and their properties – a person, his moustache, his upper lip – but we do not experience the fact that he no longer has a moustache. That is not something that we experience, but a conclusion we reach (or, in my case, fail to reach) by experiencing (perceiving) the relevant objects and their properties.

Although fact perception is, in this way, an internal affair, this does not mean it is not essentially visual. It does not mean that seeing facts (and, therefore, differences) isn’t really perception at all. Fact perception is a *coming to know by use of the senses*¹⁴ a process that ends in the head (with a perceptual judgment) but begins with the perception of those objects and properties that “reveal” this fact. Seeing that a ball is red is coming to know it is red by using your eyes, by (typically) seeing the ball (object perception) and its color (property perception) and, on this basis, making the requisite judgment: the ball is red. To see a fact, that the ball is red, then, you have to do two things: (1) see something (typically the ball and its

color), and (2) know, on the basis of what you see, that the ball is red.

Seeing a difference then, just like object and property perception, is essentially a visual phenomenon. It is a coming to know (this is the part that occurs in the head) that a difference exists by seeing (this is the part that depends on things outside the head) the things that differ. If you fail to do either (or both) of these things you fail to see a difference. I can fail to see a difference in my friend's appearance, then, in three distinct ways.

- (1) *by seeing and not knowing*: I see his moustache the first time but not the second time and fail to notice (hence do not come to believe, hence do not come to know) that there is a difference. In this case, difference blindness is not really a failure to see the relevant objects and properties, as much as a failure to use perceptually conscious information. When failing to see a difference in this way, "difference blindness" might better be described as difference *inattention* or, perhaps, difference *amnesia* (Wolfe, 1997, 1999). One is conscious of the conditions that make up the difference. One simply fails to convert this information into (conscious) judgmental form.
- (2) *by knowing and not seeing*: I find out he shaved his moustache in some way other than seeing. Perhaps I am *told* he shaved his moustache off. I don't see the difference; I am told about it.
- (3) *by neither seeing nor knowing*: I either don't see my friend or I see him after he has shaved his moustache off but I don't get a clear view of his face. I see him, but not the difference in him.

We have to be careful here. To see a difference it isn't enough that one come to know about the difference by seeing. If someone leaves me a note telling me my friend has shaved his moustache off, I come to know about a difference – a difference in my friend's appearance – by seeing, but I do not see the difference. To see a difference it is not enough to learn about it by seeing. You have to see a very particular object – presumably, the object or objects that make the difference. Or, if there is only one object involved, you have to see the properties or relations (positions, etc.) a change of which constitutes the difference. To see a difference in the position of a clock hand, I do not have to see the hand's movement (maybe its moving

too slowly for that), but I have to at least see and discriminate the visibly different positions of the hand.

If we are more inclusive, in this way, about what we count as visual perception, then seeing facts is a form of visual perception and change blindness is a genuine form of blindness. We have to remember, though, that you can be blind in this way and still see everything that is visible, everything out there in the world there is to see. Your blindness consists of a failure to do something (come to know by seeing) that is essentially visual, but your failure rests on the *coming-to-know* part of the process, the part that occurs entirely in the head, not the *by-seeing* part. This is the kind of blindness a person has who fails to see the fallacy in another person's reasoning. A failure to see something, yes, but not quite what we think of as a *visual* failure.

NOTES

¹ It isn't movement, *per se*, that attracts attention so much as the difference (perhaps local difference – see Johansson, 1950) between a target and surroundings. If everyone is waving, my friend's wave is not likely to “pop out.” We nonetheless tend to think of movement as automatically capturing our attention because we think of a “singleton” mover, an object moving in or among a group of stationary objects. There are a variety of properties (see Yantis, 1996) in addition to movement that (as singletons) capture attention in this (bottom up) way – e.g., a red spot in a crowd of green spots, a tilted line among vertical lines.

It may be of interest to note, in this regard, that Hillstrom and Yantis (1994) concluded that it is not motion (not even isolated motion) as such that captures attention. It is when the motion causes an object to segregate from its background (in the way a camouflaged animal stands out from its background as soon as it moves). Such motion, they say, creates a new perceptual *object* representation and it is *this* – the onset of a new perceptual object – that captures attention.

² Unless, of course, this arm position is salient or distinctive in some way – e.g., a raised arm in a crowd of people with lowered arms. In the text I'm imagining a scrambled scene in which there is a large variety of arm positions represented.

³ This reminds us of the mildly paradoxical fact that although change takes time, it is seen, if it is seen at all, at a time.

⁴ If, after a minute, one notices a difference in the hand's position, one can see that it has moved, but seeing that something has moved (a fact) is quite different from seeing its movement (an event). I return to the difference between seeing facts and seeing objects and events in §4.

⁵ Under normal circumstances, changes in the objects we see are accompanied by motion signals which attract a normal perceiver's attention to their location.

It is only when this local signal is swamped (via events associated with a eye saccade, a flickering image, an eyeblink, etc.) that this guidance of attention is lost and change blindness is induced (Rensink, 2000b, p. 19, 2000a, 2000c). See, also, Pashler (1998, pp. 251–252) and Rensink and O'Regan (1997, 2000).

⁶ Alva Noe (in correspondence) reminds me that there are cases of “change blindness” in which the changes are not concealed. In some experiments, for instance, an irrelevant “mud splash” occurs at the time of the change. Subjects fail to notice the change (and consequent difference) despite the occurrence of the change in full view. It is a bit like having your friend wave to you in a crowd of people all of whom are waving at you. Do you not see him wave or do you actually see his arm move and just not notice it because of the simultaneous occurrence of so many distractors?

I set aside these cases here since, as the *other* ways of demonstrating change blindness indicate, the primary interest in change blindness seems to come from an interest in our failure to detect the differences that such changes (whether seen or not) produce.

⁷ Everyone *else* (i.e., those for whom the change was not synchronized with saccades) easily saw the change and, therefore, the difference between (1) and (2).

⁸ This difference is important in evaluating what can and cannot be seen. It seems plausible to suppose that birds and dogs (not to mention human infants) see differences but, lacking sophisticated beliefs about their own perceptual states, do not believe they see differences. Unlike subjects in change blindness experiments, though, they do not believe they do not see differences. That is an important difference.

⁹ I think it is Mack and Rock's (2000) implicit acceptance of a reportability criterion of conscious awareness and, in particular, a stimulus detection criterion of conscious perception, that leads them to deny conscious perception of anything without attention (hence *Inattentional Blindness*). Things are complicated here because there are different things one can mean by attention (see Mack and Rock's discussion of this issue in Chapter 11), but I would argue (see Dretske, 1981, 1993) that what requires attention (and confers on one the consequent ability to report an awareness of the stimulus) is perception *of the fact* that there is an *x* in location *L*, *not* a perception of the *x* (*object*) in location *L*. Until we have an argument that one cannot (consciously) perceive *x* (an object) in *L* without perceiving the fact that there is an *x* in *L*, failure to detect (i.e., see the fact that there is) an *x* in *L* (accompanied by a sincere report that one has not seen *x* in *L*) will not show that one hasn't perceived (consciously) an *x* in *L*. It will only show that one hasn't seen (the fact) that there is an *x* in *L*. Although *S*'s beliefs about what he saw may be relevant evidence concerning what he saw, you can't demonstrate that *S* did not see a cufflink in the drawer by appealing to his sincere denials of having seen a cufflink in the drawer.

Merikle and Reingold (1992) have a useful discussion of different criteria for conscious awareness (perception). They say (p. 59) that many researchers feel uncomfortable measuring conscious awareness solely in terms of self-reports. A

major reason for caution is that it is difficult to know what criteria individuals are using when reporting their conscious experiences. “. . . the fundamental problem with using self-reports as a measure of awareness is that it transfers the responsibility for defining awareness from the investigator to the observer” (p. 59).

¹⁰ Given that I thought (albeit mistakenly) that my friend had gained weight, it might be said that although I saw *a* difference, I didn’t see *the* (i.e., the actual or real) difference.

¹¹ This is only “probable” because even on the present model of seeing differences (analogous to seeing spies) it is possible that subjects who say they don’t see a difference really don’t see it. People of otherwise normal eyesight who, despite looking at a spy, say they don’t see a spy, may be speaking the truth. It is possible that at the critical moment they are struck blind or suffer a *hemianopia* (lesion to the primary visual cortex that obliterates parts of the visual field).

¹² See note 11.

¹³ Shadows don’t reflect or emit light. Neither do events – games, movements, battles.

¹⁴ In a widely used and respected textbook Palmer (1999) defines visual perception as the process of acquiring knowledge about environmental objects and events by extracting information from the light they emit or reflect (p. 5). This goes too far (we can perceive objects without acquiring knowledge about them), but it is, I think, a reasonable description of *fact* perception. By tying the acquired knowledge to extraction of information from light, it makes fact perception essentially visual.

REFERENCES

- Akins, K. (1996): *Perception*, Oxford University Press.
- Bornstein, R. and Pittman, T. (eds.) (1992): *Perception without Awareness*, London: Guilford Press.
- Coltheart, V. (1999): *Fleeting Memories*, Cambridge, MA: MIT Press.
- Dretske, F. (1981): *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.
- Dretske, F. (1993): ‘Conscious Experience’, *Mind* 102.406, 1–21.
- Grimes, J. (1996): ‘On the Failure to Detect Changes in Scenes across Saccades’, in Akins (1996), pp. 89–110.
- Hillstrom, A.P. and Yantis, S. (1994): ‘Visual Motion and Attentional Capture’, *Perception and Psychophysics* 55, 399–411.
- James, W. (1950): *The Principles of Psychology*, New York: Dover (original published in 1890).
- Johansson, G. (1950): *Configuration in Event Perception*, Uppsala: Almqvist and Wiksell.
- Mack, A. and Rock, I. (2000): *Inattentional Blindness*, Cambridge, MA: MIT Press.

- McConkie, G.W. and Zola, D. (1979): 'Is Visual Information Integrated across Successive Fixations in Reading?', *Perception and Psychophysics* 25, 221–224.
- Merikle, P.M. and Reingold, E.M. (1992): 'Measuring Unconscious Perceptual Processes', in Bornstein and Pittman (1992), pp. 55–80.
- Palmer, S. (1999): *Vision Science: Photons to Phenomenology*, Cambridge, MA: MIT Press/A Bradford Book.
- Pashler, H.E. (1998): *The Psychology of Attention*, Cambridge, MA: MIT Press.
- Rensink, R.A. (2000a): 'Seeing, Sensing, and Scrutinizing', *Vision Research* 40, xxx–xxx.
- Rensink, R.A. (2000b): 'The Dynamic Representation of Scenes', *Visual Cognition* 7(1/2/3), 17–42.
- Rensink, R.A. (2000c): 'On the Failure to Detect Changes in Scenes across Brief Interruptions', *Visual Cognition* 7 (1/2/3), 127–145.
- Rensink, R.A., O'Regan, J.K. and Clark, J.J. (1997): 'To See or Not to See: The Need for Attention to Perceive Changes in Scenes', *Psychological Science* 8, 368–373.
- Rensink, R.A., O'Regan, J.K. and Clark, J.J. (2000): 'On the Failure to Detect Changes in Scenes across Brief Interruptions', *Visual Cognition* 7, 127–145.
- Simons, D.J. and Levin, D.T. (1997): 'Change Blindness', *Trends in Cognitive Science* 1, 261–267.
- Wolfe, J.M. (1997): 'In the Blink of the Mind's Eye', *Nature* 387, 756–757.
- Wolfe, J.M. (1999): 'Inattentional Amnesia', in Coldheart (1999), pp. 71–94.
- Yantis, S. (1996): 'Attentional Capture in Vision', in A. Kramer, M.G.H. Coles and G. Logan (eds.), *Converging Operations in the Study of Visual Selective Attention* (pp. 45–76), Washington, D.C.; American Psychological Association.

Department of Philosophy
Duke University
201 West Duke Building
Durham, NC 27708-0743
USA
E-mail: dretske@acpub.duke.edu