

## Fred Dretske and the Invisible Gorilla

Quassim Cassam, University of Warwick, UK

### *Abstract*

This paper is about Dretske's notion of simple seeing and his distinction between seeing in this sense and noticing. I examine various ways in which the phenomenon of inattentional blindness appears to raise questions about Dretske's account and consider how Dretske might respond to these questions. I focus on three Dretskean arguments in support of the view that in certain problem cases objects that weren't noticed were nevertheless seen. I suggest that each of these arguments is flawed even if there is a lot to be said for Dretske's attempt to account for conscious experience in epistemological terms.

1

There is a BBC radio programme called *Desert Island Discs* in which guests are asked to choose eight pieces of music they would want to have with them if they were stranded on a desert island. One can imagine a less popular variation called *Desert Island Philosophy Books*. Which eight philosophy books would you want to have with you on a desert island? I don't know that reading philosophy books would necessarily be at the top of my list of priorities on a desert island but, if I had to choose, Dretske's 1969 masterpiece *Seeing and Knowing* would be one of my eight. I have fond memories of spending time with Fred at a conference in Oslo in 2011, talking to him about his book. We didn't discuss the issues I want to raise here, but I don't doubt that he would have dealt with them with his customary charm and philosophical insight.

Early on in *Seeing and Knowing* Dretske poses a question which was still exercising him nearly forty years later, in an important 2007 paper on change blindness.<sup>1</sup> This is the question:

1

Why do we say to people, as we sometimes do, ‘But you *must* have seen it’?... It is certainly not that the person already believes he saw it (whatever *it* happens to be). Quite the contrary; he does not believe he saw it, cannot remember having seen it, or is simply dubious of whether he saw anything at all of the kind in question..... ‘You must have seen that cuff link; you were staring right at it.’ Whatever response this allegation may prompt, it is not refuted by an appeal to ignorance: ‘I did not notice it’ or ‘The drawer looked empty to me.’ He may have seen the cuff link without noticing it; he may have seen it without its looking to him *as though* there was something in the drawer (1969: 18).

The seeing at issue in this example, which I’ll call CUFF LINKS, is what Dretske calls ‘non-epistemic seeing’. In a later paper he renames it ‘simple seeing’. Simple seeing is ‘the seeing of *objects* and *things* – not *facts* about these things’ (2000a: 98). Part of what makes simple seeing non-epistemic is that you can see an object X in this sense of ‘see’ ‘*with* or *without* a knowledge of (or belief about) what one is seeing. Seeing X, just like stepping on X, is compatible with no beliefs about X’ (2000a: 101). The contrast is with *epistemic seeing*. The grammatical test for epistemic seeing is that ‘the verb “to see” has a factive nominal (that....) or a question word clause (who..., whether..., what...) as its complement’ (2000: 98a). This type of seeing is *not* compatible with no beliefs about what is seen. ‘S sees *that* X is F’ entails ‘S knows that X is F’ and ‘S believes that X is F’.

Dretske’s point about CUFF LINKS is that the fact that the person neither knows nor believes that he saw the cuff links doesn’t prove that he didn’t (simply) see them. To suppose otherwise would be to make ‘the mistake of supposing that the seer is in some special, privileged position for determining *what* is seen and, therefore, *whether* something is seen’ (2000a: 102). Furthermore, just as you can see an object without knowing or believing that

you see it, so it is possible to see an object without *noticing* or *recognizing* it. Here is another of Dretske's examples, which I will refer to as MARKETPLACE:

You are looking for a friend in a crowded marketplace. You can't find him. Later, after you've found him, 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 broad daylight. You certainly didn't see through him. You didn't, for instance, 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, blocking your view of them. So you must have seen him (2007: 217).

Just to be clear, when Dretske says 'you must have seen him', he isn't using 'see' in a sense that doesn't imply consciousness. He insists that 'seeing, hearing, and smelling x are ways of being conscious of x' (2000b: 115). So the claim is that you can be *conscious* of x without noticing or recognizing x.

What about the role of attention in simple seeing? Can you see an object without attending to it? Yes, according to Dretske, and the examples he uses to show that you can see an object without attending to it are the similar to the ones he uses to show that you can see an object without noticing it. Here is another one ('FLAG'):

Look at a flag for a few moments. How many stars do you see? If the flag has fifty stars, and none of them are obscured by folds in the cloth or by other objects, it seems reasonable to say you saw them all. *All fifty*. Which one or ones did you *notice*? It sounds odd (at least to my ear) to say that you noticed every star on the flag, but not at all odd to say you saw them all. What makes our visual experience the rich and profuse thing we know it to be is that we can see *more* than we can ever notice or attend to (2000a: 110).

To sum up, Dretske's claim is that if you see an object X, and your seeing is simple rather than epistemic, then none of the following claims is *guaranteed* to be correct: (a) you have a belief about X (b) you know something about X (c) you know that you are seeing X (d) you believe that you are seeing X (e) you notice X (f) you recognize X (g) you are attending to X. To say that these claims aren't *guaranteed* to be correct is not to say that they *couldn't* be correct. For example, you *can* notice an object you simply see, or have a belief about it. The point is that you don't *have* to notice it or have a belief about it in order to see it.

What then are the positive conditions on non-epistemic seeing? I've said a lot about what non-epistemic seeing *doesn't* require, but what *does* it require? So far, the only positive suggestion is that if you simply see an object you must be conscious of it. Dretske's positive account of simple seeing is what he calls an 'information-theoretical' account (2000a: 109). The basic idea is that 'seeing X is getting information (coded in a certain way) about X' (ibid.). Here is the same idea in a bit more detail:

We get information about things by seeing, hearing, smelling and touching them. To say that someone has *seen* X is to say that information about X has been delivered in a particular form.... When we see X, X (or some event associated with the presence of X) initiates a sequence of events that culminates in a distinctive sort of experience, the sort we call a *visual* experience. Typically, this experience embodies information about the color, shape, size, position, and movement of X (2000a: 108-9).

One can get information about X without being conscious of X, as happens in blindsight, but information in the absence of consciousness is insufficient for seeing. However, neither the receipt of information from X nor consciousness of X requires the satisfaction of conditions (a) - (g); the informational-theoretical account does not make simple seeing epistemic.

The attractions of Dretske's account are obvious. The distinction between the seeing of objects and the seeing of facts is simple and intuitive, as is the idea that objects can be seen without being noticed, recognized or known. On the other hand, any readers of Dretske who are familiar with the phenomenon of 'inattentional blindness' might worry that his account fails to account for this phenomenon. Specifically, the worry might be that seeing X, noticing X and attending to X can't be separated as sharply as Dretske claims, and that 'simple' seeing can't be quite as non-epistemic as one might suppose. My plan is this: in this section I'll describe the phenomenon of inattentional blindness and identify various ways in which it puts pressure on Dretske's account. In section 3, I'll consider some Dretskean responses to the objection from inattentional blindness. Lastly, I'll argue that these responses are inconclusive and that neither side in the debate has a wholly convincing response to the other.

I turn now to inattentional blindness. In their classic paper 'Gorillas in our midst', Daniel Simons and Christopher Chabris argue that 'there is no conscious perception without attention' and that 'when attention is diverted to another object or task, observers often fail to perceive an unexpected object, even if it appears at fixation' (1999: 1060). This is what they call 'inattentional blindness'. There are plenty of mundane illustrations of this phenomenon. You are searching for an empty seat in a crowded movie theatre and the next day your friends ask why you ignored them at the theatre: 'they were waving at you, and you looked right at them but did not see them' (1999: 1059). However, the most famous example of inattentional blindness is the invisible gorilla experiment, in which people engaged in counting the number of passes made by a group of basketball players failed to see or notice a woman wearing a full-body gorilla suit who 'walked into the scene, stopped in the middle of the players, faced the camera, thumped her chest, and then walked off, spending about nine seconds onscreen' (2011: 6). I say 'failed to see or notice' because Simons and Chabris use 'see' and 'notice' interchangeably. Here is a representative passage from their book *The Invisible Gorilla*:

Amazingly, roughly half the subjects in our study did not notice the gorilla! Since then the experiment has been repeated many times, under different conditions, with diverse audiences, and in multiple countries, but the results are always the same: About half the people failed to see the gorilla.... What made the gorilla invisible? This error of perception results from a lack of attention to an unexpected object, so it goes by the scientific name “inattentional blindness”. This name distinguishes it from forms of blindness resulting from a damaged visual system: here, people don’t see the gorilla, but not because of a problem with their eyes. When people devote their attention to a particular area or aspect of their visual world, they tend not to notice unexpected objects, even when those unexpected objects are salient, potentially important, and appear right where they are looking (2011: 6-7).

Simons and Chabris add that it isn’t just the fact of inattentional blindness that is interesting but what they call the *illusion of attention*: ‘we experience far less of our visual world than we think we do’ (2011: 7). In GORILLA and many similar cases we suppose that we should or would have seen the object in question and are surprised when we realize what we missed. This surprise is one manifestation of the illusion of attention.

Why are Simons and Chabris so confident that in GORILLA the subjects failed to see the gorilla? Aren’t they naively assuming that the seer is in some special, privileged position for determining *what* is seen and, therefore, *whether* something is seen? In MARKETPLACE when you say to your friend that you didn’t see him Dretske demurs. Wrong, he says. Given the facts as described, you probably *did* see your friend even if you didn’t notice him. By the same token, it’s easy to imagine Dretske protesting that those who claim not to have seen the gorilla are also mistaken: ‘Wrong! They probably did see it, given that it was clearly visible to them. It’s true that they didn’t *notice* it, but that is a different matter’. Conceiving of the subjects in GORILLA seeing gorilla without realizing it is no harder than, say, conceiving of

you seeing my grandmother without realizing it. Neither scenario is difficult to imagine if, as Dretske insists, people aren't authoritative about what they see or even about whether they see anything.<sup>2</sup>

The first thing to say about this argument is that the grandmother analogy is a poor one. The sense in which you can see my grandmother without realizing it is that you have a 'mistaken idea of who my grandmother is' (2000a: 102). In other words, you fail to identify the person you see as my grandmother. GORILLA is quite different. If the people described by Simons and Chabris are wrong about what they saw it isn't because they failed to identify the gorilla as a gorilla; they don't have a mistaken idea of what a gorilla is. Nor is GORILLA like another case in which people are wrong about what they can see. This is the case in which you believe you saw *something* out of the corner of your eye but are mistaken because in reality there wasn't anything there for you to see. GORILLA is not a case in which you believe you saw something that wasn't there to be seen. It is a case in which you believe you *didn't* see something that *was* there to be seen. Your reason for thinking you didn't see the gorilla is that you weren't *conscious* of it. This, you might think, is something about which you are authoritative, and if you weren't conscious of the gorilla then you didn't see it.<sup>3</sup>

This raises another question: why is it being assumed that the subjects in GORILLA weren't conscious of the gorilla? Couldn't they have been having conscious experiences they no longer remember having, or weren't conscious of having? The first of these possibilities can't be ruled out but the empirical basis for treating GORILLA as a case of memory failure isn't strong.<sup>4</sup> As for the suggestion that a person can have conscious experiences they weren't conscious of having, this sounds contradictory. Dretske argues that it isn't contradictory and that one can indeed 'have conscious experiences without being conscious that one is having them' (2000b: 125). Even so, we need a positive reason for thinking that this is the best account of GORILLA. There might be a case for dealing with GORILLA along these lines if

we are satisfied that those who denied seeing the gorilla must have been conscious of it but that is what is at issue: *why* must they have been conscious of it? Isn't this the illusion of attention all over again?

The point of these questions is not to suggest that Simons and Chabris have no case to answer. There are legitimate questions about their account, and Dretske's discussion brings these questions into focus. For example, Simons and Chabris suggest that some subjects in their experiment missed the gorilla because of a lack of attention to an unexpected object. But how does this explain CUFF LINKS and MARKETPLACE? In these cases the object isn't unexpected. You look in the drawer for your cuff links because you expect to find them there. You look for your friend in the crowded marketplace because that is where you expect him to be. You don't miss them because your attention is focused on another task. You are looking for them but still miss them. You don't notice them but why should we think that a failure to notice is a failure to see? Dretske might not be entitled to assume that you must have seen your friend or your cuff links but Simons and Chabris aren't entitled to assume that you didn't see them simply because you didn't notice them. Since 'see' and 'notice' are different concepts they need to be much more careful about moving from 'roughly half the subjects in our study did not notice the gorilla' to 'about half the people failed to see the gorilla'. Even if both statements are true, the truth of the second statement doesn't follow from the truth of the first.

Although these criticisms carry some weight they still don't address the fundamental question. No doubt Simons and Chabris aren't as careful as they might be in the passage quoted above from their 2011 book, and make argumentative moves that seem dubious from a Dretskean perspective. However, none of this affects their fundamental insight: there clearly are occasions when subjects failed to see things they think they should and would have seen, and where this failure results from inattentive blindness. Regardless of whether

the move from ‘half the subjects in our study did not notice the gorilla’ to ‘half the people failed to see the gorilla’ is legitimate, the question remains whether the gorilla was in fact seen by those who did not report seeing it. It’s not that people can’t be wrong about this kind of thing but if we refuse to take people at their word in GORILLA we need very good reasons for doing so. Just repeating that they *must* have seen the gorilla isn’t good enough. The question remains: *why* must they have seen it, and what is the *evidence* that they saw it? Similarly, in CUFF LINKS, why must the person have seen the cuff link and what is the evidence that they saw it? The challenge for Dretske is to give non-question-begging answers to these questions, that is, answers that don’t beg fundamental questions against those who are impressed by the pervasiveness of inattention blindness and illusions of attention. The next section is about whether he has the resources to respond to this challenge.

### 3

In examples like CUFF LINKS and MARKETPLACE Dretske produces a range of arguments in support of his assertion that unnoticed objects were nevertheless seen. I want to discuss three such arguments, which I’ll refer to as the Conditions Argument, the Obstruction Argument, and the Epistemological Argument. Of these arguments, the first is both the most straightforward and the least convincing. The clearest statement of it is in this comment on why we say to people ‘But you must have seen it’:

Generally speaking, we say such things in the face of a person’s disbelief: we say it when we are convinced that, despite what the person *thought* he saw, or whether he thought he saw anything at all, the physical and physiological conditions were such that the object must have looked some way to him (1969: 18).

This is from Dretske’s discussion of CUFF LINKS, and he makes much the same point about MARKETPLACE: you must have seen your friend because your eyes were open and he was standing in your line of sight only a few yards away in broad daylight. The implication is that

(i) there are physical and physiological enabling conditions for seeing (ii) in both cases these conditions were satisfied, and (iii) if the physical and physiological enabling conditions for seeing were satisfied then you must have seen the object. This is the Conditions Argument.

Unfortunately there isn't much to be said for this argument since the physical and physiological enabling conditions of perception are necessary rather than sufficient. So the fact that these conditions are satisfied doesn't guarantee that the object in question was seen. As M. G. F. Martin observes, Dretske's conditions 'determine what can be experienced, not what is experienced' (1992: 749). What is needed is mental or experiential evidence in favour of thinking that the cuff link was seen, and Dretske doesn't provide such evidence. What sort of mental or experiential evidence does Martin have in mind? He gives the example of Archie looking for a cufflink, failing to notice the cuff link in the drawer, and later thinking back to his search of the room. Having relatively good visual memory he recalls how things looked as he searched and suddenly realizes that the cuff link was in the drawer. Assuming that 'one's memory experiences typically derive from one's past perceptions' (1992: 750), his memory experience is mental or experiential evidence of how things looked to Archie when he carried out his search.

But what if no such evidence is available? Does the evidence that Archie saw the cuff link have to derive from 'within the realm of the mental' (Martin 1992: 749)? Suppose that Archie has poor visual memory and doesn't recall how things looked as he searched. He says he doesn't know where his cuff link is but gets it right when asked to guess where it might be. Wouldn't this be non-mental evidence that Archie saw the cuff link? Hardly. Even if we suppose that Archie is guessing rather than inferring the location of his cuff link, and that his guesses are highly reliable, this would at best justify the supposition that he *registered* his cuff link or was in receipt of information about it. It would not justify the supposition that he

saw or was conscious of it. The obvious analogy is with 'blindsight' rather than seeing, where people respond to visual stimuli without consciously perceiving them.<sup>5</sup>

To sum up, in CUFF LINKS Dretske tries to settle the question whether the cuff link was seen by reference to what might be called 'objective' considerations, such as the fact that certain physical and physiological enabling for seeing were satisfied in the case as described. However, such considerations are inconclusive. Since seeing X is being conscious of X what is required is experiential evidence that Archie saw the cuff link. In the absence of any such evidence, the insistence that Archie must have seen the cuff link sounds more like an expression of irritation at his having missed something so obvious than an accurate statement of what he actually saw. 'You must have seen it' is, in these contexts, another way of asking 'How could you not have seen it?'

For another example of an attempt to settle whether something was seen by reference to objective rather than mental or experiential considerations we need look no further than the Obstruction Argument. The clearest statement of this argument is in MARKETPLACE: you couldn't and didn't see the apples directly behind your friend because he was obstructing your view of them, and if he was obstructing your view of them then you must have seen him. Someone might try a similar line with GORILLA. There was a point at which you couldn't and didn't see one of the players because the gorilla was obstructing your view of her. If the gorilla was obstructing your view of that player then, whether you realize it or not, you must have seen the gorilla. This, in essence, is the Obstruction Argument.

This argument relies on the principle is that if X is obstructing your view of Y then you must see X. Is this principle correct? You might naively think that when you raise your hand in front of your face and close your eyes you can't see your hand because your eyelids are in the way.<sup>6</sup> Yet you don't see your eyelids. So this is a case in which X (your eyelid) is obstructing your view of Y (your hand) even though you don't see X. This counterexample

might be resisted on the grounds that when your eyes are closed you do in fact see the insides of your eyelids.<sup>7</sup> Nevertheless, there remains a question about the principle to which the Obstruction Argument appeals. In MARKETPLACE, the fact that your friend obstructs your view of the apples that certainly doesn't prove that you saw your friend. Again, what is missing is any experiential evidence that you saw him. For all that the Obstruction Argument says, it might be that you were getting visual information about your friend but still weren't conscious of him. If he was blocking your view of the apples maybe it follows that you or your visual system have been registering his presence but that is not the same as seeing him, being conscious of him.

Further questions about the principle to which the Obstruction Argument appeals are raised by GORILLA. Suppose that Valentina is the player who, according to the Obstruction Argument, you couldn't see because the gorilla was in the way. The Obstruction Argument concludes that you must have seen the gorilla but those who are impressed by the possibility of inattention blindness might run the argument in reverse: you don't see Valentina because the gorilla is in the way, you don't see the gorilla, so it's false that if X is obstructing your view of Y then you must see X. From a Dretskean perspective this looks question-begging, but so does the Obstruction Argument from the perspective of someone who is convinced that Simons and Chabris are right about the invisible gorilla. What is needed is some independent rationale for insisting that you saw the gorilla or, in MARKETPLACE, that you saw your friend. The best that Dretske can do is to point out that the object in question was 'standing in your line of sight only a few yards away in broad daylight' (2007: 217) but this takes us back to the Conditions Argument: the fact that something was standing in your line of sight only a few yards away in broad daylight shows that you were *in a position* to see it. It doesn't show that you *did* see it. You could be wrong about what you saw but the Obstruction Argument doesn't demonstrate that you are in fact wrong.

Dretske's best argument in support of the view that Archie must have seen the cuff link or that you must have seen your friend in the marketplace is the Epistemological Argument. This argument addresses some of the concerns I've been expressing about the Conditions and Obstruction Arguments and also casts some light on why Dretske prefers in his later writings to talk about 'simple seeing' rather than 'non-epistemic seeing'. Change blindness is the failure to see visible objects or properties whose presence or absence constitutes a change in sequentially observed scenes. As Dretske observes, 'the evidence used to support this charge of blindness is the failure of subjects to notice or detect the difference that these objects and properties make' (2007: 215). Dretske argues that not seeing the difference is consistent with seeing whatever it is that makes the difference: change blindness exhibits a failure to notice, not a failure to see.

Here is an example ('SARAH'):

Sarah looks at seven people gathered around a table. Each person is clearly visible. Sarah looks for a few seconds, runs her eye, as we say, over and, therefore foveates each person at the table but she pays no particular attention to any of them. She looks away. While looking away, an additional person – call him Sam- joins the group. Sam is clearly visible. There are now eight visible people. When Sarah looks back, she doesn't notice the difference. Having no reason to suspect a change has occurred, she thinks she is looking at the same group of people. When asked whether she sees a difference, Sarah says "No" (2007: 216).

What does Sarah see? She believes that when she looks back she sees the same people she saw the first time. Is she right? This is Dretske's verdict:

Given the situation as I've described it, it seems entirely reasonable to say that Sarah not only saw Sam the second time, but that her experience of him was of the same kind, a fully conscious experience, as was her experience of the other seven people.

This is.... object-awareness (of Sam, a person who constitutes a visible difference) without fact-awareness that there is a difference. Sarah's failure to see the difference (that there is a change) in the group of people around the table is perfectly consistent with her seeing Sam in a fully conscious way (207: 216).

What this shows, according to Dretske, is the irrelevance of change blindness to conscious perception of objects: 'change blindness is a cognitive, not a visual impairment' (2007: 217).

Why is Dretske so confident that Sarah's experience of Sam is conscious? Intuitively, the fact that Sam has just joined the group shouldn't prevent Sarah from seeing him. Why should it? However, Dretske's main argument in support of his verdict on SARAH turns on the idea that conscious experience is an 'epistemologically enabling condition' (2007: 222). In normal circumstances visual experience justifies a host of beliefs about the things we see, and Dretske's claim is that this is what makes our visual experiences conscious. For example, I know that the man I'm talking to isn't naked because I can see that he isn't: 'what makes my experience of the person conscious is that the information embodied in the experience is the basis, the justification, for my current conscious and expressible belief about the person' (2007: 219). SARAH is slightly different. She doesn't have the conscious expressible belief that Sam isn't naked but 'she would have noticed him if he had been naked or standing on his head' (2007: 219). We can even imagine Sarah being asked to testify about Sam's behavior at a wild party. Her testimony goes as follows: 'if, as you tell me, Sam was standing at that table the second time I looked, and if, as you assure me, no one was blocking my view of him, I can assure you he was not naked. None of them were. I could see that much' (2007: 219). According to Dretske, the fact that Sarah would have noticed Sam if he had been naked 'suggests' that she is getting sensory information about him.

This leads Dretske to propose what he calls an 'epistemic test for awareness' (ETA), a knowledge test for what a subject sees or is conscious of. In its simplest form this test says:

(ETA) If you can see (and thereby know) that x is F for some value of F, you must be aware of x. You can't know, by seeing, that x is F without being aware of x.

Since visual awareness of x enables *knowledge* of x it is no longer appropriate to describe it as 'non-epistemic'; 'simple' seeing is a more appropriate label. The knowledge that is at issue can be delayed, as it is in SARAH, but must be direct, meaning that your knowledge that x is F must be the result of the way x itself looks, not some information-carrying intermediary. Sarah's knowledge is delayed in the sense that she didn't think about whether Sam was naked until asked to testify. Yet delayed perceptual knowledge is 'such a pervasive phenomenon' that we often give ourselves knowledge at the time at which the experience occurred. In this sense, Sarah knew all along that Sam wasn't naked 'despite not thinking about it until asked' (2007: 220).<sup>8</sup>

The epistemic test for awareness suggests that following defence of the claim that in CUFF LINK Archie must have seen the cuff link. *Ex hypothesi* the cuff link in the drawer did not attract Archie's attention. However, it could and would have attracted his attention if we imagine circumstances to have been different. For example, he would have noticed the cuff link if it had been glowing. If we take Archie to know by seeing that it wasn't glowing then, according to ETA, he must have seen the cuff link. This is the Epistemological Argument in action, and the considerations to which it appeals are broadly mental rather than physical or physiological. If knowing is a state of mind then it is Archie's state of mind, his presumed knowledge that the cuff link wasn't glowing, that is the basis of the verdict that he must have seen it.<sup>9</sup> Of course he denies that he saw the cuff link but perhaps he would change his tune if he reflected on what he does or can know about it. In a more bullish frame of mind he might even be prepared to testify that the cuff link wasn't glowing: 'if, as you assure me, the cuff link was in the drawer, I can assure you that it wasn't glowing. I could see that much'.

Although the Epistemological Argument is in some ways an improvement on the other two arguments I have considered it is still a long way from settling matters in Dretske's favour. The first problem with the argument is one that Dretske recognizes. Going back to the case of Sarah, suppose it's true Sarah would have noticed Sam if he had been naked. Clearly this doesn't demonstrate that Sarah saw Sam. Perhaps Sarah would have noticed, and would therefore have seen, Sam if he had been naked but it doesn't follow that she actually saw him given that he *wasn't* naked and she *didn't* notice him. Consider this analogy: 'the fact that you would notice x, an object in peripheral vision, if it moved, for instance, doesn't show that you saw x when it wasn't moving. Maybe its movement causes you to redirect your gaze so that you now see it when you didn't see it before' (2007: 229). Nevertheless, the fact that Sarah would have noticed naked Sam is taken by Dretske to *suggest* that she saw him.

Even this is debatable. The alternative is to say that if Sam had been naked his nudity would have caused Sarah to notice and so to see him but there is no suggestion that she saw him when he was fully clothed. If the cuff link had been glowing Archie would have noticed and so seen it but there is no suggestion that he saw the cuff link given that it wasn't glowing. At one point Dretske comments that 'being naked in a crowd of clothed people (or, for that matter, being clothed in a crowd of naked people) affects how much attention you will attract, how easily you will be noticed, but it does not make you (as opposed to the parts of your body) visible' (2007: 220-1). That is so, but something that makes you more noticeable might also cause you to be seen, while its absence might leave you both unnoticed and unseen. The issue is not whether Sam was visible to Sarah but whether she actually saw him, and it's hard to see why we would be justified in concluding that he saw him on the basis that she *would* have seen him in the unlikely event that he took all his clothes off before entering the room.

It isn't worth spending too much time on this issue because there is a more serious and obvious problem with the Epistemological Argument. The argument takes it for granted

that Sarah would have noticed Sam if he had been naked and that Archie would have noticed the cuff link if it had been glowing. GORILLA calls this assumption into question. To see why, imagine a variant on GORILLA in which someone dressed like everyone else on screen - call him Sam- enters the scene, thumps his chest and exits a few seconds later. Sarah is shown the scene but doesn't notice Sam. Did she nevertheless see Sam? Someone who hasn't heard of GORILLA might be tempted to argue as follows: 'Sarah would have noticed Sam if he had been wearing a gorilla suit. She knows that Sam wasn't wearing a gorilla suit and if she knows that much then she must have seen Sam'. GORILLA shows clearly what is wrong with this argument: given our capacity to miss just about anything, there is no guarantee that Sarah would have seen Sam if he had been wearing a gorilla suit. She *doesn't* know that he wasn't wearing a gorilla suit, so ETA leaves it open whether she was aware of Sam.

Certainly anyone who takes seriously the illusion of attention is likely to regard what Dretske says about examples like SARAH as demonstrating his susceptibility to the illusion. The extent to which he takes us to be visually aware of our surroundings suggests that he underestimates the extent of inattention blindness, and this concern isn't alleviated by the observation that a failure to notice needn't amount to a failure to see. Sarah might have seen more than she noticed at the time but the basis of Dretske's conception of what she saw is a contentious conception of what she would have noticed in different circumstances. If, as at least some of the evidence suggests, this conception is flawed then Dretske's argument fails. It is, of course, open to him to question this evidence, but on what grounds? It might be in dispute whether those subjects in GORILLA who didn't *notice* the gorilla nevertheless *saw* it. It isn't in dispute that they didn't notice the gorilla, and this is strong *prima facie* evidence against the assumption in the variant on GORILLA that Sarah would have noticed Sam if he had been wearing a gorilla suit. By the same token, it calls into question the assumption that Sarah would have noticed Sam if he had been naked or standing on his head. For Sarah to

know that Sam wasn't naked she would need to be able to rule out the possibility that he was naked, but examples like GORILLA suggest that she is in no position to rule out this possibility.

A further worry about SARAH concerns Dretske's assumption that Sarah's supposed delayed knowledge of Sam's state of dress is a form of *perceptual* knowledge. Even if we set aside doubts about whether Sarah knows what she thinks she knows, doubts remain about the basis of her knowledge. As Dretske acknowledges, the attribution of perceptual knowledge in such cases looks charitable given the absence of any belief or judgement at the time. If, at the time Sam joined the group, Sarah neither judged nor believed that Sam was F it looks as though she didn't know that he was F. When she subsequently judges that he was F it is on the basis of what she assumes she would have noticed. In effect, she *reasons* her way to the conclusion that he was F. For example, she might reason: if Sam was wearing a clown suit I would have noticed, I didn't notice Sam wearing a clown suit, so he wasn't wearing a clown suit. Whatever one makes of this quality of this line of reasoning there is no getting away from the fact that it is an inference, and the knowledge it delivers (if any) is inferential rather than perceptual.<sup>10</sup> Sarah infers rather than saw that Sam was F, in which case ETA provides no basis for concluding that Sarah must have seen him. If there is an argument about whether Sarah saw Sam it is going to be question-begging to assume that she could *see* that he was F for some value of F.

For all these reasons, the Epistemological Argument is at best inconclusive. The issue is whether, in examples like CUFF LINKS and SARAH, we are justified in insisting that an unnoticed object must have been seen by a subject who denies seeing it. The point is not that people can't be wrong about what they saw but the insistence that they saw things they claim not to have seen must be justified. The Conditions, Obstruction and Epistemological Arguments identify a range of considerations that might be relevant, but these considerations

turn out to be much less forceful than Dretske supposes. For all that Dretske says, the question whether the unnoticed objects must have been seen remains open.

4

Although there are respects in which Dretske's position is problematic it also needs to be acknowledged that it contains insights of lasting value and significance for epistemology and the philosophy of perception. I'd like to end by listing and commenting on three of these insights, and the extent to which they are untouched by the various reservations I've been focusing on:

- (a) The distinction between the seeing of things and of facts is one that any serious theory needs to accommodate. Epistemologically speaking there is a fundamental difference between seeing an  $x$  that is  $F$  and seeing that  $x$  is  $F$ . Whether or not one wants to say that 'S sees that  $x$  is  $F$ ' entails 'S knows that  $x$  is  $F$ ' it does seem right that the cognitive and conceptual resources required to see that  $x$  is  $F$  are far more extensive than those required to see an  $x$  that is  $F$ . It is controversial how to draw the distinction between 'simple' and 'epistemic' seeing but it is hard to dispute the existence of some such distinction.
- (b) There is a genuine distinction between seeing and object and noticing it. We do indeed sometimes see things we don't notice, just as we sometimes see things we don't recognize or aren't attending to. Examples like FLAG show that these distinctions are far from spurious. The question I have been raising is not whether these are legitimate distinctions – they obviously are – but whether we see quite as much of what we don't notice as Dretske supposes. My claim is not that a failure to notice is *per se* a failure to see but there are cases in which Dretske is mistaken about what we see or must have seen. These include some cases of inattentional

blindness, and what leads Dretske astray is his view that objective considerations can be used to determine what is seen.

- (c) Simple seeing is a form of conscious experience and that what makes it conscious is its epistemological role. I have cast doubt on one of the uses to which Dretske puts his idea that conscious experience is an epistemological enabling condition but I don't dispute the fundamental idea. Seeing an object puts you in a position to know something about it.

The challenge in this area is to steer a course between two extremes, the tendency to suppose that we actually see much less than we do and the tendency to suppose that we see much more than we do. The views to which Dretske is responding might be guilty of the first of these errors while Dretske himself might be guilty of the second. Neither side accepts as decisive the considerations on which the other relies and it's hard not to feel that in some of the more contentious examples it just isn't clear what the subject did or didn't see. Did Sarah really see Sam but fail to notice him or did she not see him? When Archie protests that he just didn't see the cuff link we should be open to the idea that he is mistaken but, then again, how can we be certain what he saw? Who are we to say? Sure, he was in a position to see the cuff link but don't we often fail to see what we are in a position to see? Intuitions vary and decisive evidence is in very short supply. Dretske was a distinguished contributor to the philosophy of perception but this is a game whose rules are more obscure than even its very best players suppose.<sup>11</sup>

## REFERENCES

- Dretske, F. (1969), *Seeing and Knowing* (London: Routledge & Kegan Paul).
- Dretske, F. (2000a), 'Simple Seeing', in *Perception, Knowledge and Belief: Selected Essays* (Cambridge: Cambridge University Press): 97-112.
- Dretske, F. (2000b), 'Conscious Experience', *Perception, Knowledge and Belief: Selected Essays* (Cambridge: Cambridge University Press): 113-37.
- Dretske, F. (2007), 'What Change Blindness Teaches About Consciousness', *Philosophical Perspectives* 21: 215-30.
- Martin, M. G. F. (1992), 'Perception, Concepts, and Memory', *Philosophical Review* 101: 745-63.
- Simons, D. J. & Chabris, C. F. (1999), 'Gorillas in our midst: sustained inattention blindness for dynamic events', *Perception* 28: 1059-1074.
- Simons, D. J. & Chabris, C. F. (2011), *The Invisible Gorilla and other Ways our Intuitions Deceive Us* (London: HarperCollins).
- Weiskrantz, L. (1986), *Blindsight: A Case Study and Implications* (Oxford: Oxford University Press).
- Williamson, T. (2000), *Knowledge and its Limits* (Oxford: Oxford University Press).

---

<sup>1</sup> Dretske 2007.

<sup>2</sup> The grandmother example is one that Dretske uses to make the point that people aren't authorities about what they see: 'Whether or not you saw my grandmother is a question that in some situations I am in a better position to determine than you. If you, with a mistaken idea of who my grandmother is, smugly assert that you did not see my grandmother today, I can correct you by pointing out that you certainly did: she was the woman to whom you gave your seat on the bus' (2000a: 102).

<sup>3</sup> The claim that this is something about which you are authoritative might, of course, be challenged.

<sup>4</sup> At any rate, this is what Simons and Chabris argue. They maintain that 'if observers did consciously perceive and then forget the gorilla, they presumably would not be particularly surprised when asked if there had been a gorilla in the display. Yet, observers in our study were consistently surprised when they viewed the display a second time' (1999: 1072). This is one of a range of empirical arguments against the 'forgetting' hypothesis. See Simons and Chabris 1999: 1072-3.

<sup>5</sup> For an account of this phenomenon see Weiskrantz 1986.

<sup>6</sup> No doubt there is a much more sophisticated physiological story about why closing your eyes prevents you from seeing your hand.

<sup>7</sup> Thanks to Craig French for this response to the counterexample. See also this post from Eric Schwitzgebel: <http://schwitzsplinters.blogspot.co.uk/2006/10/can-you-see-insides-of-your-eyelids.html>.

<sup>8</sup> What Dretske calls 'delayed' perceptual knowledge might be regarded as inferential. It doesn't follow that it isn't perceptual, at least on the assumption that perceptual knowledge is always to some degree inferential, but the role of inference in Dretske's example harder to overlook than its role in ordinary perceptual knowledge. Strictly speaking, one might think, Sarah infers *rather than* saw that Sam wasn't naked, and she infers that he wasn't naked from the fact that she didn't notice him, together with the assumption that she would have noticed him if he had been naked. See below for more on this.

<sup>9</sup> On the idea that knowing is a state of mind see Williamson 2000.

<sup>10</sup> I don't mean to imply that 'perceptual' and 'inferential' are, in general, contraries. See note 8 above. The point is that *in this case* we seem to have a piece of knowledge that is inferential rather than perceptual.

<sup>11</sup> Thanks to Craig French for very helpful comments on an earlier draft of this paper.