King's Research Portal

DOI:

[10.1111/phpr.12166](https://doi.org/10.1111/phpr.12166)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

McNeill, W. E. S. (2015). The Visual Role of Objects' Facing Surfaces. *PHILOSOPHY AND PHENOMENOLOGICAL RESEARCH*. <https://doi.org/10.1111/phpr.12166>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

The Visual Role of Objects' Facing Surfaces

William E. S. McNeill

University of York

ABSTRACT

It is often assumed that when we see common opaque objects in standard light this is in virtue of seeing their facing surfaces. Here I argue that we should reject that claim. Either we don't see objects' facing surfaces, or – if we hold on to the claim that we do see such things – it is at least not *in virtue* of seeing them that we see common opaque objects. I end by showing how this conclusion squares both with our intuitions and the facts of vision science.

I

We see common-or-garden objects: apples, birds and icebergs for instance. And we see *parts* of those objects: the iceberg's tip; the bird's wing. Moreover as Jackson put it, 'we often see things in virtue of seeing other things' (1977: 19). I might count as seeing the iceberg *in virtue of* seeing its tip; as seeing the bird *in virtue of* seeing its wing. In general 'I am properly said to see an opaque object if I see a reasonably substantial part of it' (1977: 19).¹ Consider the following example:

Iceberg You're on the deck of a ship looking out to sea. You see a white, triangular object in the distance. It is the tip of an iceberg. In virtue of seeing this, you see the iceberg.

Much of the iceberg remains unseen, hidden as it is beneath the waterline. The iceberg's tip is not identical with the iceberg itself. They differ both in spatial properties and persistence conditions, for example. Nonetheless you count as seeing the iceberg. And you do so in virtue of seeing much less than the iceberg. This is so even if you fail to grasp the significance of what you've seen.

No opaque object is such that at any particular time *all* of it can be seen. After all 'we can't see the far side and the inside' (Clarke 1965: 99).² But many opaque

¹ Read literally – as offering a sufficient condition – this sounds right. It is harder to identify *necessary* conditions – see e.g. Dretske 1969 pp27ff; Warnock 1955; McLaughlin 1984.

² I will neglect the possibility of complex arrangements of mirrors. These might deliver to one the synchronic seeing of all of an opaque object's surface. But they won't deliver its innards. They will deliver only a jumbled,

objects are nonetheless seen. So it is tempting to generalize; to claim that ‘any such object is seen in virtue of seeing some [proper] part of that object’ (Sorensen 2008: 20).³

Suppose the object we see is partly occluded – the iceberg by the waterline, the bird by some branch, say. It is natural to suggest that the part of these objects we see, in virtue of which we count as seeing the objects themselves, is that part of them which is unoccluded: the tip or the wing. But now consider a slightly different case:

Apple You’re looking at your clean white desk in good light. On the desk is an apple. You see the apple.

You do not see the apple’s innards or its far side. But at the same time your view of it is entirely unoccluded. Which part of the apple do you see, in virtue of which you count as seeing the apple *per se*? We will have to say that the part of it you see is its facing surface or ‘front layer’ (Sorensen 2008: 48). What the apple’s facing surface would occlude is the apple’s *remainder*. As Armstrong put it, ‘an apple is seen in virtue of seeing much less than the apple’ (1979: 136).

In this way we are led to the view that ‘we see opaque objects by seeing their surfaces’ (Sorensen 2008: 44). Restricting ourselves to frontlit conditions,⁴ the part of those surfaces we see, in virtue of which we see common opaque objects, is whatever part of their facing surface remains unoccluded. Where the object is in plain view, we see it in virtue of seeing its facing surface. Call this the *Surface View*.

The *Surface View* should be distinguished from one on which seeing an object’s facing surface *blocks* you from seeing the object of which it is a surface. On the *Surface View* the fact to be explained is how you count as seeing apples, birds, trees and so on themselves. In line with the *Surface View* I will assume that, in standard conditions, you can see common objects such as apples.

A surface in the relevant sense is a physical object. It is a proper part of the object of which it is a surface; a *layer*. This contrasts with an object’s boundary.⁵ A boundary is the interface between an object and its environment. It is not a part of the object, nor is it a part of the object’s environment. It has no depth. It does not interact with light.

A *facing* surface is that physical part of the object which faces you. So long as your view of the object is unoccluded, its facing surface is that part of the object

distorted, overlapping or disconnected view of its surface. And these kinds of complications won’t be relevant here.

³ Cf. Neta 2007, p.61

⁴ I exclude Sorensen’s silhouettes for now. The argument has repercussions for silhouettes, however. I discuss this in section 4 below.

⁵ For more on the distinction see Stroll 1988, especially chp 3.

from which light will, in standard front-lit conditions, stimulate your retina. In such conditions it is ‘where the action is’ (Sorensen 2011: 207).

Putting things this way brings out a causal motivation for accepting the *Surface View*. It is only the facing surface of the apple which is having a causal impact on you. If the apple’s remainder were discretely vapourized the visual scene would be indistinguishable. This might suggest to you that it is the apple’s facing surface which is the basic object of perception – that, if you see the apple, this is in virtue of seeing its causally active part; its facing surface.

There are complications here. Geometrical and causal extensions of ‘facing surface’ may differ, for instance when an object is travelling near the speed of light (Sorensen 2011). And we can see silhouettes, where the causal ‘action’ is at the absorption layer of the silhouetted objects’ *rear* surfaces.⁶ For the purposes of this paper I will focus on frontlit objects travelling at normal speeds, though I discuss silhouettes briefly in section 4.

You can see the apple or the iceberg without seeing them *as* apples and icebergs –or indeed, *as* anything in particular. The claim is extensional. And you can see an apple or an iceberg (I will assume) without noticing them, attending to them, or being able to recall your visual contact with them. The sense of ‘see’ at issue here is *simple, non-epistemic, or transparent*.⁷

The *Surface View* is that you secure this visual relation to these kinds of things by virtue of seeing other kinds of things: their unoccluded facing surfaces. I will take it that the claim that you’ve seen such a surface is also extensional. You needn’t attend to the apple’s facing surface or see it *as such*. You needn’t notice it or be able to recall your visual contact with it.

For these reasons we should distinguish the *Surface View* from the claim that you see objects in virtue of experiencing certain of their features. Rather, it is the view that it is by being perceptually related to one kind of physical object in your environment – surfaces – that you count as being perceptually related to other kinds of objects in your environment – apples, icebergs and so on.

Stroll takes the visibility of Venus and Jupiter to undermine the assumption that all visible, opaque objects *have* surfaces. If Stroll were right then it could at least not always be by seeing opaque objects’ facing surfaces that you see the opaque objects themselves. His arguments are simple. (1) We can see Venus. But the surface of Venus is obscured by a thick, impenetrable cloud of hot acid. Hence while we see Venus we do not see Venus’ facing surface. (1988: 75ff) (2) Jupiter is a gas giant. It has no surface. (1988: 31) So while we see Jupiter there

⁶ Cf. Sorensen 2008, chp. 1. I discuss silhouettes further in section 4. For the purposes of this paper I will leave aside discussion of the perception of very fast objects.

⁷ The first two labels come from Dretske (1969; 1981), the third comes from Jackson (1977).

is no facing surface in virtue of which we could see Jupiter. In a similar vein Austin protests ‘what and where exactly is the surface of a cat?’ (1964: 100).⁸

The causal motivation for the *Surface View* presents us with an answer, and with it a solution to the objection. Perhaps common parlance doesn’t ascribe to all visible, opaque objects neat, solid, visible surfaces. Nonetheless at any time when an opaque object is seen, a part of it will be causally active in stimulating the viewer’s retina, while most of it will not. That part may be structurally very complex – the fluffy facing side of the cat. It may be very thin, like the reflective surface of a snooker ball. But it could be hundreds or thousands of miles thick. Light penetrates deep into the gas clouds of Jupiter and some of that light returns to stimulate the astronomer’s retina.

To individuate an object’s facing surface the vision scientist goes as deep as is needed to find the causal action but no deeper (Sorensen 2008: 46). In this way we can always identify a relevant surface. In frontlit conditions where an opaque object is seen there is always a potential referent for ‘facing surface’.⁹ In frontlit conditions, we can *define* the facing surface as that part of the object where the relevant causal action is.

If the *Surface View* is correct then so too are the following:

- a) You see the apple’s facing surface
- b) It is in virtue of seeing the apple’s facing surface that you count as seeing the apple *per se*

If either of these turns out to be incorrect then so too will the *Surface View*. In what follows I will not argue unconditionally for the falsity of either. Instead I will argue for the following disjunction:¹⁰

Either

- i) you do not see the apple’s facing surface

or

- ii) you see the apple’s facing surface, but it is not in virtue of seeing this that you count as seeing the apple *per se*.

If the disjunction is true then the *Surface View* is not. And the truth of the disjunction turns on two ways of understanding what would be involved in

⁸ Stroll also discusses at some length more subtle examples, where the perceiver’s interests affect their willingness to assent to the claim that they’ve seen an object’s surface. I discuss this further in section 4.

⁹ Contrast Neta 2007, p. 62. The solution is at the mercy of our understanding (or lack of it) of the nature of causation. I do not have space here to do these further questions any justice. You can think of it as an assumption for the sake of argument. It provides a base on which to construct the *Surface View*.

¹⁰ This should be distinguished from perceptual *disjunctivism*. My discussion here is neutral as between disjunctive and intentional accounts of perceptual experience. On either theory, you will in non-hallucinatory cases be perceptually related to some objects in your environment. The questions are then *how* and *which ones*?

seeing the apple's facing surface. On one understanding you don't meet the conditions for seeing the apple's facing surface, hence (a) is false. That result is highly counter-intuitive. There may yet be a weaker sense in which you see the apple's facing surface. But on this understanding we lose the explanatory force implicit in (b) – that seeing the apple's facing surface *explains the sense in which* you see the apple *per se*.

In the next section I assume that you could count as seeing the apple's facing surface only if you visually differentiated it. I argue that you cannot visually differentiate the apple's facing surface, so could not count as seeing it. In section III I treat the arguments in section II as reason to give up the assumption that in order to see the apple's facing surface you have to visually differentiate it. I argue that only if you visually differentiate the apple's facing surface could seeing it be that in virtue of which you see the apple. Whether or not you see the facing surface, it follows from the fact that you do not visually differentiate it that the *Surface View* is false.

Nonetheless it may still feel contrary to good sense to deny that we see objects in virtue of seeing their facing surfaces. In section IV I try to allay some potential doubts. I argue that the conclusion is compatible with treating objects' facing surfaces as playing a crucial visual role, in line with both scientific fact and our intuitions.

II

For now accept Dretske's well-cited thought that for any subject *S*, *S* sees *O* only if *O* "is visually differentiated from its immediate environment by *S*" (Dretske 1969: 20). The question I address in this section is whether – if we accept this condition – it remains true that you see the apple's facing surface.

Dretske's claim invites the question of what it is to visually differentiate an object. And it is hard to give a neat positive account of what this involves. The sense of 'see' at issue here is simple or non-epistemic. So in order for *S* to visually differentiate *O* it is not necessary that *S* find *O* salient or that *S* attend to or notice *O*. *S* need form no particular beliefs about *O*. Nonetheless I will take the thought to be that in some minimal sense, *O* must visually *stand out* to *S*. Reflect on cases where you *fail* to see *O*.

Simple Sticker You're looking at a matte red wall in good light. Carefully placed in the centre of the wall is a thin matte red sticker of the same hue as the wall. From where you stand you don't see the sticker. The wall provides the perfect camouflage.

It would only be by getting up close that you would see the sticker – only then that its edges would give it away.¹¹ Or consider ‘a white moth on a white tree trunk. The white moth causes part of what I see although not in a way that enables me to discern its outline. I am looking at the moth but I am not seeing it’ (Sorensen 2008: 54).

In both cases an object in front of you fails – in even a minimal sense – to *stand out* against its backdrop. In neither case do you see that object. The objects don’t look any way such that they are visually distinct from what surrounds them. Camouflaging an object often involves making it blend in with its environment in this way.¹²

The suggestion is that having an object visually stand out to you is a necessary condition on your seeing that object. As Campbell put it ‘it is not just having the mountain in my visual field that matters. I have to single the thing out visually, I have to see it as a figure against a background’ (2002: 7).¹³

The apple you see *stands out* against its backdrop. After all, *there it is*. The problem is that if this is right, the apple’s facing surface does not stand out against *its* backdrop.

The apple and its facing surface differ in their mass, dimensions and persistence conditions. They are distinct objects. So they are two different possible objects of perception; two different things that you might see.

While the apple is an object distinct from its facing surface, it is in its immediate vicinity. So the facing surface’s backdrop includes the apple. But the apple’s backdrop does not include the apple. So the apple’s backdrop differs from that of its facing surface.

Along with proponents of the Surface View I am assuming that you see the apple. So the apple is not only an aspect of its facing surface’s backdrop but a part of its immediate visible environment. In this way the apple differs from its remainder, which you do not see.

We have, as it were, two distinct objects ‘stacked up’ in the same portion of your visual field. And supposing for a moment that we see both, these two objects cannot differ in their look. They would take up the same space in the visual field, would bear the same visual contrasts with their environment. They would be visually indistinguishable.

¹¹ Cf. Dretske 1969: 23. In his example the colour is *beige*.

¹² It doesn’t always. See for example Dretske’s humanoid Martians (1973). Martians camouflage themselves by looking just like humans. They remain perfectly visible. They just make sure they are not *identifiable*.

¹³ Cf. e.g. Siegel 2006: 434. It may be that Campbell has in mind here a less minimal understanding of what it would take for an object to ‘stand out’ to you. The question will turn on how we interpret the ‘as’ in ‘see it as a figure against a background’.

The red sticker is visually indistinguishable from its red wall backdrop. This is why you don't get to visually differentiate the red sticker. If the apple's facing surface is visually indistinguishable from its apple backdrop this should mean that you don't get to visually differentiate the apple's facing surface. It blends in visually with the apple. You *just see the apple*.

It is true that the apple's facing surface is visually distinct from the apple's surroundings. While this does not distinguish the apple's facing surface from everything in its immediate environment, we know that all visual differentiation of opaque objects must be partial. For example, you don't visually differentiate the apple from what the apple occludes.

Coin You're looking at the apple on your desk in good light. There is a coin directly behind the apple. The apple is not visually differentiated from the coin it occludes. Nonetheless you visually differentiate the apple.

In *Apple* the facing surface is no less visually differentiated from what surrounds the apple than the apple itself. And just as in *Coin*, in order to visually differentiate the apple's facing surface it needn't stand out as distinct from the apple's remainder. So for the same reasons that you see the apple in *Coin*, you might be tempted to conclude that you see its facing surface.

However while you see the apple in *Coin*, you do not see the apple-coin complex. The apple-coin complex is not a part of the visible environment that could interfere with your seeing the apple *per se*.¹⁴ Conversely while all participants agree that you cannot see the apple's remainder, all agree that you *can* nonetheless see the apple *per se* – a surface-remainder complex.

So the apple is relevant to the question of whether you visually differentiate its facing surface in a way that neither the coin, the apple-coin complex nor the apple's remainder are. None of these three objects are seen. While they are all parts of the facing surface's environment, they are not parts of the facing surface's *visible* environment.

Nonetheless the apple's facing surface is still partially differentiated from its immediate environment at its edge. To see how an object whose edge is partially differentiated might nonetheless remain unseen reflect on an adaptation of *Sticker*.

Protruding Sticker Imagine that the right-hand side of the wall before you is painted red. The left-hand side of the wall is painted

¹⁴ Though see Neta 2007. Neta constructs an example where we would naturally say that you see a wall-ladder complex, where the wall fully occludes your view of the ladder. Neta accepts that it takes special circumstances to generate this effect – in this case the fact that the wall and ladder comprise an artwork. The implications of such facts for the current debate are moot.

white. At the intersection is placed a large, round, matte red sticker of the same hue as the wall's right-hand side.

The red sticker stands out relative to the white left-hand side of the wall. I take it that you still fail to visually differentiate *the sticker*; that *this object* doesn't visually stand out to you. It *blends in* to the painted right-hand side of the wall.

As in the original *Sticker*, the protruding sticker here alters how things look to you. It makes the red right-hand side of the wall look as if it has a large dimple protruding into the wall's white left-hand side. And as in the original *Sticker*, the protruding sticker is not seen by you. The example differs from the original in that you do visually differentiate what is the left-hand edge of the sticker. Nonetheless while looking at the sticker you're not visually differentiating the sticker itself from *enough* of its immediate surroundings to count as seeing *it*.

The protruding sticker contrasts with some of its immediate environment – the wall's white left-hand side. Likewise, the apple's facing surface contrasts with some of *its* immediate environment – the white table. Nonetheless the sticker is not visually differentiated because of the interference of another aspect of its immediate, visible environment – the wall's red right-hand side. Likewise the apple's facing surface is not visually differentiated because of the interference of another aspect of *its* immediate, visible environment – *the apple*. Just as the red sticker blends in visually with the wall's red right-hand side, the apple's facing surface blends in visually with the apple of which it is a surface.

The apple and its facing surface are distinct. But it would be arbitrary to distinguish them on visual grounds. *If* we see the apple, and *if* seeing an object requires visually distinguishing it from its immediate visible environment, then we do not see the object which is the apple's facing surface.

Contrast this with the case of the iceberg. The iceberg and the iceberg's tip are distinct objects. And it is *not* visually arbitrary to distinguish them. It is not that the iceberg and its tip *look different*. You see the iceberg in virtue of seeing its tip and you do not see the iceberg's remainder. There is only one visual achievement involved. Nonetheless the tip is a distinctive visual presence in a way that the apple's facing surface cannot be. The iceberg's tip bears visual contrasts with its environment distinct from those of the iceberg. It takes up a distinctive space in the visual field.

Much of the iceberg is *visibly* occluded, by the waterline surrounding the tip's base. Conversely, you cannot visibly discern the boundary between the apple's facing surface and its remainder. And while the iceberg's tip is visibly bound at its lower edge by the waterline, the waterline is no part of the iceberg's boundary. Conversely the apple and its facing surface inevitably share their visual boundaries.

If you drag the iceberg out of the water you see more iceberg. But there's nothing you can do to see more apple. You can see more *of* the apple by turning it round. By turning the apple around you see more of its features. In this way you can get a better look at the apple. But your original look at the apple was entirely unoccluded. You're perceptually related to no more *mass* of apple when you turn it round. By dragging the iceberg out of the water you get a better look at the iceberg. But you've also got yourself perceptually related to a greater *mass* of iceberg by doing this.

We can imagine peeling away the apple's facing surface and discarding the remainder. In this case you do see the apple's facing surface. But this is just a different case, with a different object there in front of you. We can imagine peeling off the apple's facing surface and moving it slightly towards you so that the apple's remainder is still there, occluded by this separated facing side.¹⁵ You now see (just) that facing surface. But again this is simply a different case. Where the apple occludes the coin you see just the apple. You see neither the coin nor the apple-coin complex. Where the apple's facing surface and its remainder are unattached, you see just the facing surface. You see neither its remainder nor the surface-remainder complex – the apple itself. But where the undoctored apple is still there in front of you, you *do* see the apple itself.

In visual terms, the relation between the iceberg and its tip is quite unlike the relationship between the apple and its facing surface. To distinguish between the apple and its facing surface is visually arbitrary. You see the apple. Yet its facing surface cannot be a visually differentiated part of the apple that you see. Given all this it would follow from accepting Dretske's condition as here understood that you don't see the apple's facing surface. You just *see the apple*.

III

For this among other reasons we may want to abandon the visual differentiation condition, at least in its full generality. Dretske himself notes a 'limiting case': you have your nose pressed against a large white wall. You see the wall. But it is not clear how you could visually differentiate the wall from its immediate environment. 'In this position it has no environment, and so one could hardly be expected to differentiate it from one' (Dretske 1969: 26). In a slightly different case, if the two halves of a ping pong ball are placed one over each eye you see each half from the inside before the ganzfeld effect takes hold (cf. Sorensen 2008: 244).¹⁶

Moreover it is true of both the right- and left-hand sides of the red wall in *Sticker* that you see them (Dretske 1969: 25). But neither side of the wall visually 'stands out' in the intuitive way that motivated the visual differentiation

¹⁵ Cf. Thompson Clarke's (1965) example of the pen and Neta 2007: 55ff.

¹⁶ Thanks to an anonymous referee at this journal for bringing this example to my attention.

condition. If you see the wall's right-hand side but you don't visually differentiate that side of the wall, then visual differentiation is not a general condition on object seeing.¹⁷

In his discussion of undifferentiated seeing Dretske focuses particularly on surfaces. *S* 'runs his eye' over a large expanse of wall. Someone then traces a triangular shape on the wall and asks whether *S* saw this part of it:

we can say that *S* saw that section of the wall simply on the condition that *that part of the wall* looked some way to him at the time he saw the wall. Differentiation is unnecessary when the individuality of what is seen rests *solely* on such arbitrary tracing operations or *exclusively* on its relative spatial location and configuration. (*ibid.*)

What *S* is said to have seen in this undifferentiated way is a part of the wall's *surface*. If surfaces are a special case this will be relevant in deciding whether (or in what sense) you've seen the apple's facing surface.

So for one reason or another we might want to hold on to the idea that you see the apple's facing surface while conceding that you don't visually differentiate it from its immediate surroundings; that it doesn't *stand out* for you as distinct from the apple itself.

The problem for the *Surface View* now is that it is unclear in what sense it would be in virtue of seeing its visually undifferentiated facing surface that you would count as seeing the apple *per se*.

It doesn't follow from the claim that you see both the apple and the apple's facing surface that you see the former in virtue of seeing the latter. You see the brown patch towards the top of the apple. But you do not see the apple in virtue of seeing that patch, nor do you see the patch in virtue of seeing the apple.

To say that you see *B* in virtue of seeing *A* is to say that your seeing *A* somehow explains or *makes sense of* the fact that you see *B*. Reflect on the iceberg case again. You count as seeing the iceberg in virtue of seeing its tip. Securing your visual relation to the iceberg's tip explains how it is true that you see the iceberg *per se* on this occasion.

The form of explanation at work here may or may not be thought causal. You might think that seeing the iceberg's tip *causes* you to see the iceberg itself. Analogously, claiming that you took a Eurostar train would explain how it is that you find yourself in Paris by identifying what caused you to be there.

As an interpretation of 'in virtue of' this is disputed. Jackson argued that the relation is not causal.¹⁸ Your car is red in virtue of its body being red. But the

¹⁷ See French *forthcoming* for further candidate examples.

¹⁸ See also Scott Campbell's (2004) discussion on the 'in virtue of' relation.

redness of the car body does not *cause* the car to be red – ‘the spray painting does that’ (1977:16). Rather, Jackson thinks that we should treat the relation as manifesting the analysis of one fact in terms of another. We analyse the fact that the car is red in terms of the fact that the car’s body is red but not vice versa. As he puts it ‘there is an asymmetry in the notion which it is difficult to capture’ (1977:17). As I will put this, while the fact that the car is red entails that the car’s body is red, it does not explain it. Explanation here runs in the other direction.

One way or the other, the claim that you’ve seen the iceberg’s tip is explanatory. It answers the question of how or in what sense it is true that you’ve seen the iceberg on this occasion.

Whether causal or not, the form of explanation at work here is not counterfactually robust. It does not follow from your failing to see the iceberg’s tip that you’ve failed to see the iceberg. You might be a diver seeing it from underneath. The claim that you’ve seen the iceberg’s tip isolates a condition on seeing the iceberg which given the context is sufficient for seeing the iceberg and, unlike the alternatives, is actual. While the Eurostar caused you to be in Paris it could have been Air France. While the car touches the kerb in virtue of its front left tyre touching the kerb it could have been the rear right.

However hold all else equal. If from the deck of your ship you lose visual contact with the iceberg’s tip you lose visual contact with the iceberg. If from your vantage point the iceberg’s tip is not visually differentiated from its immediate surroundings you don’t see the iceberg. Seeing the iceberg’s tip can play the explanatory role it does only because the iceberg’s tip visually stands out to you.

Suppose you do see the apple’s facing surface. You still don’t visually differentiate it from the apple *per se*. If you do not visually differentiate the apple’s facing surface then the claim that you’ve seen it cannot play the kind of explanatory role that seeing the iceberg’s tip plays.

Another way to bring out the disanalogy between *Iceberg* and *Apple* here is to reflect on the extent to which object seeing is context-sensitive. ‘Do we see Harold when we only see the tip of his finger?’ (Dretske 1969: 27). ‘The conversational context and one’s particular interests ... [will] affect the question of whether enough of *D* was seen to see *D*’ (*ibid.*: 28).¹⁹

Even if we accept that ‘seeing a relevant attached part of an object suffices for seeing the object’ (Sorensen 2008: 44), which attached parts will count as ‘relevant’ may alter in line with the perceiver’s context and perhaps also the context of attribution. ‘In general, seeing *x* requires not merely seeing a part of *x*, but seeing a sufficient and characteristic part of *x*. But how much is sufficient,

¹⁹ Cf. Warnock 1955; McLaughlin 1984; Neta 2007

and which parts are characteristic?’ (Neta 2007: 61). Neta argues that ‘the answers to these questions depend in part ... upon the context in which the seeing ascription is made’ (*ibid.*).

But note that in Dretske’s example it is not in doubt that we’ve seen Harold’s fingertip. It follows that the question of whether or not we’ve seen Harold is sensitive to context in a way that seeing his fingertip is not.²⁰

In general, in cases of part-whole seeing, the question of whether you’ve seen the part (tip, wing, or fingertip) is less open to context than the question of whether you’ve seen the object of which it is a part (iceberg, bird or Harold). After all Harold and the iceberg may be partly occluded but their tips are not.

We should treat this as manifesting the direction of explanation I’ve tried to bring out in the discussion so far. Your visual relation to the iceberg’s tip is not subject to the vagaries of context in the way that your relation to the iceberg is. In other words, your visual relation to the iceberg’s tip is more solid or basic than your visual relation to the iceberg itself. This is one reason why seeing the tip explains how it is that you’ve seen the iceberg and not vice versa.

So, if we see the apple by virtue of seeing its facing surface we would expect the question of whether we’ve seen the apple, in a particular situation, to be more open to context than the question of whether we’ve seen its facing surface in that same situation.

The question is whether this is what we find. And in spite of some outward appearances to the contrary I suggest that it is not. Perhaps you can be enticed to withdraw from the *claim* that you see the apple. As Austin put it, you might not always be prepared to *stick your neck out* (1964: 99). Yet we know that what we see and what we can reasonably claim to see sometimes come apart. Indeed this is a key motivation for isolating a notion of ‘simple’ seeing.

A pink elephant wanders into the field before you. You have good reason to think that you’re hallucinating. It would be odd for you to assent to the claim that you see the pink elephant – much more reasonable to ask for a doctor. Nonetheless you see the pink elephant. While the understandable belief that you’re hallucinating defeats your justification for believing that you’re seeing a pink elephant it is not in a position to undermine your visual relation to the unusual creature standing directly in front of you. However well justified, the belief that you’re hallucinating doesn’t make you go blind.

Now reflect on Clarke’s (1965) ‘surface enquiry’. You start by agreeing that you’ve seen the apple (or Clarke’s *cheese*). But a philosopher keeps pressing you

²⁰ Though I doubt it is, I’ve phrased this in a way designed to leave it open that seeing the fingertip is open to context in some way.

on how much of the apple you've seen until you begin to agree that it could only really be its facing surface – after all, the rest of it is out of sight.

For Clarke, the result of this enquiry is that your perceptual situation has altered. Where before you were visually related to the apple now you are visually related to the apple's facing surface. So for Clarke altering the conversational context can affect whether or not you see the apple itself.

To accept this would amount to denying that there is any space for the notion of 'simple' seeing. Given that there can be no genuinely *perceptual* distinction between the apple and its facing surface the change that the surface enquiry affects can only be one of attention or belief. If your attention or your beliefs affect which unoccluded objects you see then an assumption I share with proponents of the *Surface View* is lost.

So long as we leave room for the notion of 'simple' seeing, attending to the puzzle about how much of common opaque objects you see cannot undermine your visual relation to unoccluded objects sitting right there, in good light before you. You visually differentiate both the pink elephant and the unoccluded apple, whatever your worries. Both 'stand out' in the minimal sense I pointed to above.

So at the very least, the question of whether you've seen the apple's facing surface is no *less* context-sensitive than the question of whether you've seen the apple itself.

And the question of whether you've seen the apple is not context sensitive in the way that the questions of whether you've seen Harold or the iceberg are. Why not? In contrast with your views of Harold or the iceberg, your view of the apple is entirely unoccluded.

It makes sense to say that you see the iceberg in virtue of seeing its tip. So there is nothing wrong with 'in virtue of' claims in general. But in *Apple*, the claim that you see its facing surface cannot play the explanatory role that the tip played in *Iceberg*. There is nothing perceptual that could distinguish your seeing the apple's facing surface from your seeing the apple *per se*. So even if we are persuaded to hold on to the claim that you see the apple's facing surface the idea that you see the apple *in virtue* of this becomes empty.

If anything it is tempting to say that explanatory direction is reversed in this case. In *Iceberg* the visual achievement is to differentiate the iceberg's tip. In the context you get the iceberg for free. In *Apple* the object you visually differentiate is the apple itself. If you get its undifferentiated facing surface at all, you get it for free. One way or another, even if you see the apple's facing surface it is not in virtue of seeing this that you count as seeing the apple.

At this point we might be tempted to isolate an explanatory asymmetry between the apple and its facing surface in the following way. Whenever you see the

facing surface of an apple you see the apple. But it doesn't follow from the fact that you've seen the apple that you've seen *that bit* of it. You could have seen it from the other side. So pointing to the apple's facing surface explains how you see what you do in a way that merely pointing to the apple could not.

The sense in which seeing the apple's facing surface is explanatory here would be thinner than the sense I've invoked so far. But it is plausible that it would be enough to rescue the claim that it is in virtue of seeing the apple's facing surface that you see the apple. And that is all a proponent of the *Surface View* was committed to.

But the defence doesn't work. Like 'horizon', 'facing surface' is indexical. In frontlit conditions, from whichever direction you view the apple you would always see its facing surface. So it would always follow from the fact that you'd seen the apple that you'd have seen its facing surface.

In fact the asymmetry that this defence invokes is not objectual but *featural*. Pointing to a particular part of the apple's surface helps explain which features of the apple you've seen. If you'd seen it from the other side you may have seen the bruising. The mere claim that you've seen the apple leaves it open whether you've seen this feature of it or not.

But the *Surface View* was the claim that you see the apple in virtue of seeing an *object* distinct from the apple. It was not a claim about what you see *of* the apple. To bring out the contrast note that you can see an object while seeing *none* of its features.²¹ You're colourblind so don't see the redness of the wall. You're wearing distorting spectacles so don't see its straight edges or right-angled corners. You're at one end of an Ames room so don't see its true location. Nonetheless you see the wall. The wall's features play a crucial causal role both in getting you visually related to the wall and in fixing how things look to you – but not in virtue of your *seeing* them. I return to this theme below.

IV

To sum up: you do not visually differentiate opaque objects' facing surfaces – visually, they blend in with the objects of which they are surfaces. Were it true that you could count as seeing a facing surface only if you visually differentiated it then you do not see facing surfaces. Yet if we accept that you can see facing surfaces without visually differentiating them, it cannot be in virtue of this undifferentiated contact that you see the objects of which they are surfaces. On either disjunct the *Surface View* is false.

In this Section I tackle two questions: How does the conclusion relate to our intuitions? And how does it relate to the empirical facts of vision science?

²¹ Cf. Tye 2009, p95 on cases of ubiquitous error.

It is highly counter-intuitive to suggest that we *don't see* opaque objects' facing surfaces. The argument I have set out is compatible with our seeing such facing surfaces in a visually undifferentiated way. But it is instructive to start by considering that initial intuition.

The claim that we don't see opaque objects' facing surfaces may be counter-intuitive for the wrong reasons. Stroll notes examples where it is quite natural to say that we *don't* see an object's facing surface, even when we see the object of which it is a surface. A pilot flying high enough might naturally say that she cannot see the surface of the lake below.²² If an object is just a speck on the horizon, it doesn't sound unnatural to say that you can't see its surface. The same would be true if you were to view a snooker ball spinning very fast in front of you. Our linguistic intuitions here vary with context.

At the same time the flexibility of such intuitions relies on our interest in objects' surface *features*. The pilot cannot 'make out' the lake's surface; cannot tell whether it is choppy or calm. You cannot 'make out' the snooker ball's surface; cannot tell whether it is pitted or smooth, patterned or not. Our willingness to hold on to the claim that we see objects' facing surfaces appears to wax and wane in proportion to our contact with those objects' surface features. Stroll puts things more boldly: 'we must be able to discern some detail in the surface of an object before we can be said to be seeing the surface' (1988:86).

So if in general we assent to the claim that we see objects' facing surfaces this may be because, in general, we take ourselves to be able to make out some of their surface features.

Seeing a standard object's surface features is compatible with failing to see that object which is its facing surface. If our commitment to the claim that we see objects' facing surfaces derives from an intuition about our relation to objects' surface *features*, it doesn't give us any direct motivation to hold on to the claim that we (non-epistemically, transparently or 'simply') see facing surfaces.

To put the thought into relief, consider our natural response when it comes to silhouettes. Intuitively when you see the solar eclipse you see the silhouetted moon itself and not simply its shadow or its edge. But to return to the motivations for the *Surface View*, which part of the moon do you see in this backlit situation? It cannot be its causally inert facing side. It must be its causally active absorption layer or rear surface. 'Since the front surface of a silhouetted object is idle, the causal theory rules out the possibility that we see silhouettes by seeing their front surfaces. Only back surfaces are available' (Sorensen 2008: 44). In this way we are led to the counter-intuitive conclusion

²² Cf. Stroll 1987, p392

that when we see the silhouetted moon we see it in virtue of seeing its *rear* surface.

These thoughts about silhouettes present us with two potential lessons. One is that more theoretical considerations may sometimes trump more intuitive ones. Sorensen *bites the bullet*. A defender of the view that the visual differentiation condition should hold for facing surfaces may also decide to bite the bullet. The failure of facing surfaces to meet the visual differentiation condition could be considered a reason to deny that we see facing surfaces – in spite of our intuitions. As an aside – and I return to this a little later – note that doing so would provide a suprisingly simple and intuitive answer to the puzzle about silhouettes. We *see the (silhouetted) moon* and not just its effects or its edge; yet we do *not* see its rear surface (phew!).

Another potential lesson would be that perhaps there is not much of a bullet to bite here. A defender of the visual differentiation condition with respect to facing surfaces can pin the intuition that we *do* see frontlit objects' facing surfaces on our interest in surface features. Sorensen can accept that we are not in receipt of a rich seam of information about the moon's rear surface. He can pin the intuition that we *don't* see that object on the same featural interests.

So: we can't make use of our natural allegiance to the claim that we see objects' facing surfaces to support the *Surface View*. But at the same time we cannot use my diagnosis of the motivation behind that willingness to undermine the claim that we *do* in some sense see objects' facing surfaces. Neither our intuitive acceptance of the claim that we see facing surfaces nor my story about the featural motivation for that acceptance are decisive here.

If not on the basis of linguistic intuition, how then can the issue be decided? We are left with the two motivations I sketched for the *Surface View* in section 1. On the one hand is the thought that if we see opaque objects this must be in virtue of seeing only some *part* of them. On the other hand is the thought that in frontlit conditions only the facing surface is causally active.

To start with the latter: accept that 'in frontlit conditions, we see an object by virtue of the light transmitted by its front layer' (Sorensen 2008: 48). There is no simple move from this either to the claim that we see objects' facing surfaces or that we see objects in virtue of seeing their facing surfaces.

While being a part of the causal chain which leads to our seeing what we do may be a necessary condition on being seen, apples meet this criterion no less than their facing surfaces. Or at least if they do not then the *Surface View* is false – we wouldn't be able to see apples at all.

Being a part of the causal chain which leads to our seeing what we do is not sufficient for being seen. Light rays are part of that causal chain but are at least

very rarely things we see. What light rays do is to *enable* us to see what we do. Usually they are not a part of the visual field but a means of generating it.

The same goes for the molecules that make up the objects we see. Whether or not we see objects' facing surfaces, we surely don't see their molecular constituents. Yet it is the way light interacts with these constituents that explains how our retina come to be stimulated as they do. In short, being explanatory does not require being seen. We can account for the causal role facing surfaces play without accepting the *Surface View*.

The causal condition may be thought to fare better in association with the thought that we see opaque objects only in virtue of seeing some proper part of them. Suppose that when we see the apple this must be in virtue of our seeing some proper part of it. Where the apple is unoccluded only causal considerations can identify a relevant part. Causal considerations give us the apple's facing surface. It would follow that when we see the apple this must be in virtue of seeing its facing surface.

Surprisingly though, causal theorists who accept the *Surface View* are committed to denying that whenever you see an opaque object this is in virtue of seeing a proper part of it. The facing surface of the apple is an opaque object. According to these theorists it is an object that you see. But on pain of regress – Leonardo's 'whittling objection'²³ – you don't see the object's facing surface by seeing a proper part of *it*. In virtue of seeing what would you see *that* part? And you don't see the object's facing surface by seeing every part of it either – you don't see its molecular parts.

For one reason or another we have to allow that – at least for some objects – you can see them yet not in virtue of seeing either a proper part of them *or* all of them. Some opaque objects you *just see*. So there's no particular motivation for the *Surface View* here either. An opponent could claim that it was the apple that you *just saw*.

And note again that this claim is at least highly intuitive when we're talking about silhouettes. It is open to an opponent of the *Surface View* to claim that we *just see* the silhouetted moon, the causal intervention of its absorption layer being merely an enabling condition for this basic perceptual achievement. There's no simpler way to account for our intuitions about silhouettes. We have to reject the thought that whenever we see opaque objects this is in virtue of seeing one of their proper parts anyway. And when we do, good things happen.

There is a weaker assumption that these considerations do not undermine: the simpler claim merely that if you see an opaque object then you see some part of it.

²³ Sorensen 2008: 44. Cf. Stroll 1988: 40-46; Neta 2007: 59

It doesn't follow from your seeing some part of an object that you see the object *in virtue of* seeing that part. You do undoubtedly see some parts of the apple. You see the brown patch towards the left, the stalk, the speckles towards the top and the shaded area to the lower right, for example. While some of these are parts of the apple's facing surface, they are by the same token parts of the apple. Unlike the apple's facing surface they are parts of the apple which visually stand out.

You also see the left- and right-hand sides of the apple. These are parts of the apple that you see even though they do not visually stand out. Accept that the apple's facing surface is a part of the apple that you see in this undifferentiated way. Then we can also accept that whenever you see an opaque object you see its facing surface. It could still be false that you see opaque objects *in virtue of* seeing their facing surfaces.

So the considerations up until now are compatible with denying my first disjunct – that we don't see objects' facing surfaces. We could see them in an undifferentiated way. The second disjunct remains true. If we only see facing surfaces in an undifferentiated way then we don't see standard objects *in virtue of* seeing their facing surfaces.

Likewise we could hold on to the view that we see silhouettes *rear* surfaces in this undifferentiated way. But what would be gained? The motivation for saying that we see facing surfaces was that they look some way to us. In the same way if someone traces around the camouflaged red sticker in the simple case and asks if we can see that patch we have to say "yes". It is a visible part of the red wall, after all. It makes a 'positive contribution' (Dretske 1969: 23) to the light array that stimulates our retina. In this sense the sticker is not *invisible* even though it is not something we visually differentiate; not an object that 'stands out'.

We may have become suspicious about the way in which our desire to say we've seen things varies relative to our perceptual contact with some of their features. But the claim that an object looks some way is weaker than the claim that we can discern any of its features. It might look ways it is not. The 'red' wall might be a white wall cunningly illuminated.

So linguistic suspicions notwithstanding, we can still motivate the view that we see the camouflaged sticker or the apple's facing surface in undifferentiated ways. The lake's surface still looks *some* way to the high-flying pilot: blue, maybe dark; *featureless*, perhaps.

If you accept the *Surface View*, it follows from your seeing the silhouetted moon that you've seen its rear surface. It is in virtue of seeing some part of it that you see the moon. Neither its edge nor its shadow are eligible. And causal

considerations identify only its rear surface. This is why some are led to reject the view that we see the silhouetted moon at all.²⁴

I've put forward grounds independent of puzzles about silhouettes for rejecting the *Surface View*. If we reject the *Surface View*, it does not follow from your seeing the silhouetted moon that you see its rear surface. Our motivation for holding on to the claim that you see objects' *facing* surfaces in frontlit conditions is that they look some way to you. But does the moon's rear surface look any way to you? The *moon itself* looks blackish or silhouetted, round and unperforated.²⁵ Its rear surface, I suggest, looks no way at all. When we give up on the *Surface View* can take or leave the idea that you see some part of a particular object's surface in line with our intuitions.

In short, then: rejecting the *Surface View* does not undermine our intuitions. It accounts for more of them. And it is not incompatible with the facts of vision science. Facing surfaces underpin counterfactually robust claims about how things look to you. 'Rigid objects can be almost entirely hollowed out without disturbing their appearance' (Sorensen 2008: 46). Facing surfaces explain which features of the objects you see are visible. For the purposes of explaining such facts we can treat objects' facing surfaces along the lines of light rays or molecules.

And unlike icebergs' tips, facing surfaces are not visually differentiated. So for the purposes of explaining how we see standard objects it is no *use* to claim that we see their facing surfaces. *Seeing* the apple's facing surface can't explain seeing the apple if seeing the apple's facing surface just is, amounts to, or is constituted by seeing the apple itself.

In other words the claim that in frontlit conditions we see objects' facing surfaces is neither necessary nor sufficient to explain how we see those objects. It is unnecessary because we can treat our causal contact with facing surfaces as a condition which enables us to see what we do. It is insufficient because in frontlit conditions the distinction between an object and its facing surface is visually arbitrary.

To conclude: I have argued that you don't see common-or-garden objects in virtue of seeing their facing surfaces. This argument is compatible with the claim that their facing surfaces play a relevant causal role and with the claim that their

²⁴ It is tempting to *tollens* Sorensen's *ponens*. But the variety of ways in which this can be done shows that doing so would not be sufficient to undermine the *Surface View*. For this reason among others I've tried to undermine the *Surface View* independently of any implications it has for our seeing objects' rear surfaces. I take our ability to both accept that we see silhouetted objects and deny that we see their rear surfaces as a benign by-product of my conclusion rather than an argument in itself.

²⁵ Thanks to an anonymous reviewer at this journal for these examples. For more see Sorensen 2011.

facing surfaces look some way to you. It is compatible with the claim that whenever you see an opaque object you see some part of it and even that – in a surprisingly uninteresting way – you see opaque objects’ facing surfaces.

The apple and its facing surface differ both ontologically and in the way they contribute causally to your perceptual awareness. Nonetheless when you see an object in plain view, you do not see it in virtue of *seeing* anything less than the object itself. When the apple is right there in front of you just *see the apple*.²⁶

BIBLIOGRAPHY

- Armstrong, D. M. (1979). Perception, Sense Data and Causality. In G. F. MacDonald (Ed.), *Perception and Identity: Essays Presented to A. J. Ayer* (pp. 84-98). London: Macmillan.
- Austin, J. L., & Warnock, G. J. (1964). *Sense and sensibilia* (Vol. 83): Oxford University Press Oxford.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Campbell, S. (2004). Seeing Objects and Surfaces, and the 'In Virtue Of' Relation. *Philosophy*, 79(3), 393-402.
- Clarke, T. (1965). Seeing Surfaces and Physical Objects. In M. Black (Ed.), *Philosophy in America*.
- Dretske, F. (1969). *Seeing and Knowing*. London: Routledge.
- Dretske, F. (1973). Perception and Other Minds. *Noûs*, 7, 34-44.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, Mass.: CSLI.
- French, C. *forthcoming*. 'Object Seeing and Spatial Location'. In: F. Dorsch et al. (ed.) *Phenomenal Presence*. Oxford, Oxford University Press.
- Green, M. S. (2007). *Self-Expression*. Oxford: Oxford University Press.
- Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Mclaughlin, B. P. (1984). Perception, causation, and supervenience. *Midwest Studies in Philosophy*, 9(1), 569-592.

²⁶ Earlier versions of this paper were presented at the Institute of Philosophy, the Joint Sessions 2014 in Cambridge, and to the Reason & Mind group at the University of York. Thanks to participants for their helpful questions and comments. I was aided by discussions with Rory Madden, Craig French, Keith Allen, Paul Noordhof and Tom Stoneham. For their many helpful comments on earlier drafts of this paper, thanks are due to Craig French and an anonymous reviewer for this journal.

- Neta, R. (2007). Contextualism and a puzzle about seeing. *Philosophical Studies*, 134(1), 53-63.
- Siegel, S. (2006). How Does Visual Phenomenology Constrain Object-Seeing? *Australasian Journal of Philosophy*, 84(3), 429-441.
- Sorensen, R. (2008). *Seeing Dark Things*. Oxford: Oxford University Press.
- Sorensen, R. (2011). Silhouettes: A Reply from the Dark Side. *Acta Analytica*, 26(2), 199-211.
- Stroll, A. (1987). Seeing Surfaces. *Midwest Studies in Philosophy*, 10(1), 379-398.
- Stroll, A. (1988). *Surfaces*. Minneapolis: University of Minnesota Press.
- Tye, M. (2009). *Consciousness revisited: Materialism without phenomenal concepts*. MIT Press.