Through a Shadow, Darkly[*]

István Aranyosi

(Draft of 30 April 2014, do not cite yet, forthcoming in Mac Cumhaill & Crowther, *Perceptual Ephemera*, OUP)

I

The dictionary tells you that a shadow is a dark area or volume caused by an opaque object blocking some light. The definition is correct, but we need to clarify a couple of its elements: darkness and blocking.

Regarding darkness we should note two things. One is that a shadow does not need to be pure darkness; depending on the circumstances, a shadow can be darker or lighter in absolute terms. The second thing to note is that a shadow to exist at a certain location does not always require a contrast with lighter areas surrounding it. Such a contrast is required in order to see a shadow, but not for the shadow to exist as such. To illustrate this point, consider the shadow cast by a box on a white wall. Suppose we painted the area on the wall surrounding the shadow in a color that exactly matches the darkness of the shadow. Did we render the shadow invisible? Yes, it is not visually identifiable. But did we annihilate it? I would say we did not. As far as ontology goes, the cast shadow is simply the absence of light on a certain surface as a result of its being blocked by an opaque object, or, in other words, as a result of an opaque object standing between the light source and the surface.

In Aranyosi 2007, I accounted for the location of shadows in terms of where light would have been, had it not been blocked. Now I realize that the phrase "light would

[*] I would like to thank Clare Mac Cumhaill and Thomas Crowther for organizing (and inviting me to speak at) the conference *Perceptual Ephemera* at the University of Geneva, in 2012, where some of the ideas for this paper were conceived. I would also like to thank the *Logos* Research Group in Analytic Philosophy, in Barcelona for hosting me when I was working on this paper, and in particular (and no particular order) to Dan Zeman, Joseph Moore, José Díez, Josep Prades, Josep Macià, Dan López De Sa, and Aurélien Darbellay for feedback on my presentation of a first draft of it in November 2013. Last, but not least, I am grateful to my Bilkent collegues for excellent feedback during a work-in-progress seminar in April 2014, namely, Bill Wringe, Jack Woods, Sandrine Berges, Varol Akman, Kourken Michaelian, and Sandy Berkovski.

have been at location S" is ambiguous in a way that makes the difference between cases when an opaque object blocks light such that we do have a shadow cast on a surface and those when we don't have a shadow even if there is a blocker. When giving an account of shadows, the phrase should really mean "light would have been *reflected* at location S", rather than merely "light would have hit at location S". The reason is that an object cannot cast a shadow on a completely black surface, that is, a surface that does not reflect any wavelength of light. The darkness of the surface is not the effect of the object blocking some light, because had the blocker been removed, the surface would have been as dark as before. Compare now with our previous case—that of darkening the area around the shadow cast on the white wall. If the blocker had been removed, there would have been light reflected by the white wall. This is what it means for a shadow to exist at a location.

Let us now turn from ontology to perception. What does it mean to see a shadow? If the existence of a shadow at S means that light would be reflected (or reflected to a higher extent) by the surface at S, were it not blocked, then seeing a shadow at S appears to mean either seeing that at S light would be reflected, were it not blocked, or seeing the light that would be reflected at S, were it not blocked. Both are problematic.

The first is an instance of what Dretske (1969) calls "epistemic seeing", or “seeing that”, and it is implausible as an analysis of seeing a shadow, because it requires too much. We can sometimes see a shadow without being aware that it is a shadow, and, more generally, seeing a shadow does not require that one knows or believes that it is the result of some object blocking light. The second interpretation is an instance of seeing as such, or nonepistemic seeing, and, although it is definitely more plausible than *seeing that*, it has its own problem in our case. When we see a shadow at S, it is not the light that, counterfactually, would be there, were it not blocked, that we see. We see the darkness, the absence of light. On a causal theory of seeing, it is even be hard to make sense of a causal connection to a *counterfactual* cause of our experience, as suggested by the phrase "to see the light that would be at S ...". In fact, when our attention is focused on a shadow, it is in virtue of not seeing the light that would have been reflected at S that

we see the shadow. But if it is grounded in our not seeing the light that otherwise would be there, were it not blocked, what is it that we actually see?

II

We can sometimes gain insight into what it means to A by raising the question of what it means to fail to A. As noted above, one can fail to see a shadow at location S when in fact a shadow exists at S. Is it the lack of contrast with the area surrounding S, when the light is prevented by a blocker from reaching and being reflected at S, that makes us fail to see a shadow? Our case of darkening the area around the shadow on the white wall seems to point in that direction, but this is not always the case. The hatchet fish is a bioluminiscent ocean dweller emitting light downwards in a way that matches the properties of the sunlight from above, which the fish perceives with its eyes. As this feature was selected for its camouflaging properties, biologists talk about the hatchet fish as "camouflaging its own shadow". But, strictly speaking, the fish annihilates its own shadow; a shadow needs to be dark. So here we have a case of light being blocked by an opaque object, that is, light that would be reflected by the surface were it not blocked, and a lack of contrast generated by the light source on the fish's belly, yet, there is no shadow. So the lack of contrast with the surroundings of S is not sufficient to explain failure to see a shadow at S even when otherwise there is a blocker of light.

Is it, then, the lack of contrast with the surroundings of S and the darkness of the surface at S that together explain cases of failing to see the shadow at S? It is, but we should dwell a bit into what we mean by "darkness" in this case. We shouldn't imply that a shadow is visible only on the condition that it is darker than its surrounding surface. Consider again our example with the box casting a shadow on the white wall, and suppose we painted the wall around the shadow with pure black. The shadow would appear as much lighter than the surrounding area. We would see the shadow even though probably we would fail to recognize it as a shadow, if we weren't aware of the setting. We

might even think that the box is *illuminating* the wall, yet, we would actually see its shadow.

Darkness, therefore, should be understood here as implying a contrast not with the reflectance properties of the surface surrounding the shadow, but as a contrast with a counterfactual situation involving the surface occupied by the shadow itself. So maybe we fail to see a shadow at S when the surface at S is (1) actually visually indistinguishable from the surroundings of S and (2) would appear less dark than actually, if the light were not blocked. Condition (2) is an existential condition for shadows, stating their causally dependent nature, hence, it will have to also be satisfied when we do see a shadow.

At this point one might raise the following case as a potential counterexample to the above analysis. Suppose the box casts a shadow on the white wall, and that we paint the area surrounding the shadow the same color as the shaded area, so they become indistinguishable visually. Suppose that the material surface exactly occupying the area on which the shadow is cast is photosensitive, so that it adjusts its level of light absorption in such a way as to keep its apparent darkness constant under all illumination conditions. So the above counterfactual condition does not hold: were the blocker removed, the surface under the shadow would appear exactly as dark as before. If we agreed that there was a shadow on the wall before removing the blocker, then this looks like a counterexample to my claim about the conditions in which we fail to see a shadow.

But why should we accept the claim that there is a shadow on the wall in such conditions? The cause of why the surface is dark is not exclusively and immediately that light is blocked by an opaque object. Its darkness is caused by the intrinsic material features of the surface itself. Shadows, on the other hand, are exclusively and immediately grounded in causal processes extrinsic to the surface itself. This is not to say that the intrinsic properties of the material surface play no role, but that they are to be considered background conditions rather than active causes. Being background conditions, they are supposed to be kept constant when analyzing a causal connection

involving them in terms of counterfactuals. So I don't think this is a counterexample to the above analysis.

III

If the above characterization of failing to see a shadow at S is correct, then we see a shadow at S when the surface at S is visually distinguishable from the surroundings of S, and the surface at that location would reflect more light, were the blocker removed. Yet, even this is not sufficient. Consider your hand casting a shadow on the white wall. Now move it closer and closer to the wall until it touches its surface. The shadow did not disappear at the moment you have touched the wall since it is still true that the surface touched by your hand would have reflected more light, had your hand been removed.[1] The visual appearance of the surface of your hand also contrasts with its surroundings, namely, the white wall. Yet, you fail to see its shadow.

The explanation is that you can't see the surface through your opaque hand, but you do sometimes see it *through* a shadow, except, of course, when the shadow is very dark. To see a shadow, as reflected by the discussion so far and assuming a causal account of seeing, is to have a visual causal relation to the shaded area, rather than to the shadow itself. To see a shadow is to have your visual awareness of a surface *modulated* by the absence of light caused by a blocker, in such a way that a visual contrast with the surrounding area is also noticeable. So it is a way of failing to see a surface, or, to be more precise, a degree of not seeing a surface as well as in normal conditions of

[1] Jack Woods and Bill Wringe pointed out to me that this claim appears to entail that there are plenty of (indeed, infinitely many) shadows inside any opaque object that receives some light, because removing any proper part of such an object will reveal a surface that reflects more light than before the removal (as it was inside the opaque object, the revealed surface failed to reflect any light whatsoever). This would be a rather peculiar view to endorse. But, in fact, there is no such entailment. In my analysis I talk about *surfaces* that would have reflected more light, had light not been blocked. By definition, no interior proper part of an opaque object counts as a surface. "Surface" is understood to mean outer skin or envelope of an object. See Casati and Varzi 1994: ch. 2.

illumination.[2] The highest such degree is when a shadow is totally dark, for instance, when the blocker is close enough to the surface and the light is very strong. The lowest such degree is when you are still able to notice a shadow, as a consequence of its contrast with the surrounding (though, again, you might still not be aware that it is a shadow[3]). We cross this limit of seeing a shadow when the contrast it creates on the surface is not noticeable any longer unless, for example, the blocker moves. Sorensen offers such an example based on his personal experience:

> (...) faint shadows can disappear when they move too little. This first came to my attention while marveling at how thoroughly my lawn had been cleared of leaves. To my amazement, a leaf seemed to materialize on the bare grass and then disappear. It was actually the shadow of a leaf in a tree. The shadow was too faint to see when stationary but became visible when a breeze moved the leaf. Although the leaf shadow did this repeatedly, I was unable to break the camouflage of the shadow. (2008: 85)

The idea that seeing a shadow is a degree of failing to see a surface is not crazy. The animal world offers indirect support for it through the phenomenon of countershading, the most common body surface pigmentation phenotype in the animal kingdom. Countershading is a luminance gradient from a dark back to a light belly. We observe it in terrestrial diurnal mammals, birds, and reptiles, in amphibians, as well as in most fish species. The function of countershading is to camouflage the animal by counterbalancing the effects of self-shadowing. If the animal had a uniform coloration over its body surface, it would appear as a solid object with an attached shadow in the lower part of its

---

[2] Normal conditions of illumination will involve a range of light intensities, such that, for instance, the case in which I see a shadow but the light surrounding it is so intense that were the blocker removed, I would be blinded does not count as a counterexample to my proposed analysis.

[3] For instance, you are looking in the morning at a number of slightly darker, barely noticeable grease stains on your white shirt. One of the dark patches is actually a shadow.

body and a lighter, illuminated upper region of its body, hence easily identifiable by predators. By creating a coloration gradient that is opposite to this default coloration and luminance condition, Nature in effect has brought about a "fake shadow" on the upper region of the animals body, meant to conceal that surface. There is a counterexample to this "dark back and light belly" pattern in the class of fish, namely, one species of catfish, *Synodontis nigriventris*, which has a darker belly and lighter back. However, it is also known as the upside-down catfish, as it swims, unlike other fish species, upside-down, thus reaffirming the enormous selective pressure and advantage which lead to the emergence of countershading in the first place.

IV

Sorensen (2008: ch. 4) asks us to imagine the shadow cast by a spinning sphere. The sphere is spinning, but does its shadow spin too? Sorensen argues that it does. The reason is that the shadow is a causally derivative entity, entirely dependent on its caster; hence, it inherits the causally relevant properties of the caster, in this particular case its properties related to motion. The only reason we are hesitant in ascribing rotational motion to the shadow of the spinning sphere is that we do not see the shadow as moving. In order to persuade us of his view, Sorensen asks us to consider the case when the spinning sphere has a stick attached to it (2008: 82). Here, we seem to intuit that the shadow is indeed spinning, as the shadow of the stick is moving around the center of the sphere's shadow. Now suppose the stick falls of at some point. It would be counterintuitive to say that all of a sudden, just because the stick fell off, the sphere's shadow stopped spinning.

This intuition-pump is based on the compelling thought that we should maximize uniformity in judgment about some phenomenon, when the change in the setup of the mechanism that generates that phenomenon does not appear to be so drastic as to destroy it or radically change it. Uniformity in judgment, on the other hand, has to be based in the uniform application of some theory about the relevant phenomenon—in Sorensen's

particular example, it is the Newtonian theory of motion as applied to shadows. However, Sorensen's view that the shadow is spinning is not the only way to ensure such uniformity. If what I have been arguing for is correct, the apparent motion of shadows is, indeed, *merely apparent*, and the sought-after uniformity in theory at least, if not in judgment, is ensured by taking into account the behavior of such appearances. More to the point, motion of shadows is difference in location of the surface we fail to see properly. Speed is rate at which location of what we fail to see changes over time. Neither rotational, nor translational motion of shadows is real in the sense of being an objective property "out there"—it is a property of what, or rather *where*, we fail to see properly to a degree or other. However, apparent translational motion—which is the case of the attached stick's shadow in Sorensen's example—is indeed a case of a change of location of a surface we fail to see properly, and so unlike the rotational motion of a fixed disk or sphere. Hence, saying that the shadow has stopped moving when the stick fell off, though it does express a discontinuity in judgment, it does not amount to a discontinuity in *theory*. The uniform application of our theory requires a non-uniform judgment in this particular case.[4]

V

In the "disappearing act" puzzle, Sorensen (2008: ch. 3) asks us to visualize a completely black truncated cone brick which is a perfect imitation of the shadow cast by a conical object suspended below a lamp. We now slide the brick under the cone so that it exactly occupies the space where the shadow was/is located. Here, we either see the shadow, in which case we totally fail to see the object, or we see the black object, in which case we properly see the black surface. However, since the brick does not absorb light (it is in

[4] For an elaboration of the view that the only type of motion we should attribute to shadows is translational and deformational, see Sayan 1996.

complete darkness), it must be the case, according to Sorensen, that what we see is the shadow.

But there is no shadow, because it is not true that the object would reflect more light, were the blocker removed. We don't see the object either, because it does not absorb any light. We fail to see anything located at S, but pure darkness, hence, not all spatially delimited regions of darkness are shadows or black objects.

VI

In the "intersecting eclipses" riddle, Sorensen (2008: ch. 1) asks us to imagine looking at an eclipse that is the result of two planets, Far (farther from the observer) and Near (closer to the observer), lining up in the direction of the sunlight. One question is which of the two planets do we see, and Sorensen argues that it is Far, because it is the only of the two that is causally relevant in casting the shadow over the observer. However, Sorensen goes further and argues that when we see silhouettes, that is, backlit objects, what we see is the far ends of these, that is, the backs of these objects, because those are the surfaces that are causally relevant in blocking the light. This is an interesting view and the argument for it seems compelling. However, I think there is a better alternative.

First, *pace* Jonathan Westphal (2011), silhouettes are not shadows, or front views of shadows. Consider you own silhouette in the mirror. You see both the apparent reflection of your silhouette (see below for details), and your shadow cast on the mirror. Now paint the portion of the mirror where the shadow is located black. Your shadow has been annihilated; no shadows on black surfaces. Yet, you still see your silhouette in the mirror.

We see objects in the mirror indirectly, by seeing their image directly. This image is a reflection when the object that we see reflects light. When it doesn't, it is a *parareflection* (Sorensen 2008: ch. 6). Shadows and black objects in the mirror are such pararflections. However, silhouettes in the mirror are not even parareflections. The reason

is that seeing parareflections of shadows and black objects is still a form of indirectly seeing these particulars, but an apparent reflection of a silhouette does not qualify as indirectly seeing the silhouette. A photograph of a shadow presents to us an image of the shadow, which is not itself a shadow. A mirror parareflection of a shadow does the same thing as the photograph; it presents an image of the shadow. If the apparent reflection of the silhouette in the mirror were an indirect way to see the silhouette itself, it should not itself be a silhouette, but an image of it. But it is not an image.

To see this, note that the only reason one might think we see the silhouette indirectly through the mirror is that we see the light around it indirectly, because it is reflected by the mirror. But not all mirror reflections are images. The mirror reflection of an illuminated object is an image, but the reflection of light as such is not. When the Sun is reflected by a mirror, we see the sunlight itself, not an image of it. The same is true for any light source's mirror reflection; we see the light itself, not an image of it. Since the silhouette's ontology is determined by that of the light that surrounds it, the apparent reflection of the silhouette in the mirror is, therefore, not an image. It is the silhouette itself. But how can we see the silhouette in the mirror when in fact it is located on the opposite side, where the light blocker is located? The answer is that, just like in the previous case of Sorensen's disappearing act, we don't see something. The silhouette is pure darkness; whereas seeing a shadow is a degree of failing to see a surface or a volume, seeing a silhouette is not even that much. A particular can't be in two or more places at the same time, and so we can't directly see one at a location while it objectively inhabits another. But we can fail to see one at such a location. Seeing a silhouette translates into a particular failure to see; it is ontologically dependent on seeing the light itself that is not blocked. As long as what we directly see at different locations, S1 ... Sn is the same particular quantity of light, our failure to see something at S1... Sn is the same particular failure.

One might object as follows. What we see in the mirror is *a* silhouette. However, it is not *the* silhouette. If it were, we would have a doubly located particular (one on each side of the object on the line connecting the light source and the specular space. So what

we see is another silhouette. But a silhouette according to the objector's own understanding, just like a shadow, is always the silhouette of a concrete object, whereas the silhouette that we are supposed to see in the mirror, if the objection works, is not; there are no concrete objects in specular space.

VII

In more recent work, Sorensen (2011: 206) mentions the phenomenon of amodal completion as adding further evidence to his view that by seeing silhouettes we see the backs of the blockers, and furthermore, that we even see the geometric structure of the interiors of those back surfaces.

> "Action in a silhouetted setting brings a conviction that one is seeing in a robust fashion. We prolong our basketball game past sundown because the silhouetted net has a discernible three-dimensional structure. From a long distance, you see well enough to side step an opponent, jump, aim and throw. The other players look up, see the ball sail across the lingering light of the sky, bounce against the backboard, teeter along the hoop, and then drop through the net.
>
> As a supplement to the arguments in *Seeing Dark Things*, I now add an appeal to "amodal completion". When a dog stands behind a picket fence, only segments of the dog are visible. Despite the occlusions, we see the dog as a whole object, not as a scattered sample of dog parts. The same holds when we see a silhouetted scene of a dog behind a fence."

This appeal to amodal completion can be generalized and used as a further argument for the view that shadow perception is no different from object perception since they both share our visual system's tendency to amodally complete objectively disconnected parts

of the perceptual object. Amodal completion is the phenomenon by which our visual system groups disconnected parts of appearances in a way that matches the continuity of the concrete objects that cause our perception, although in terms of retinal stimulation these parts remain represented as discontinuous. Textbook examples include the dog behind a picket fence, which we perceive as a whole dog rather than as a collection of disconnected dog slices, or the dotted line pattern, which we perceive as a line rather than as a mere collection of disconnected line segments.

However, this further, amodal completion based potential rationale for taking shadow perception as a form of and on a par with perception of ordinary objects is not very convincing once we realize that, in fact, amodal completion of shadows is not an autonomous phenomenon, but one that is parasitic on, or an offshot of, amodal completion in the perception of ordinary objects, and that it only occurs in cases when we perceive shadows with familiar shapes, that is, with shapes of ordinary objects. We amodally complete the dog's shadow just because its shape looks like that of a dog. If we consider randomly shaped superposed objects, it is far from clear that their shadows appear to us in a way in which we amodally complete each of them. Even the case of an interrupted simple, quasi-linear shadow, analogous to the dotted line in Gestalt psychology, it is not at all clear whether we amodally complete it so as to perceive it as one shadow. An example is the optical phenomenon phenomenon known as the *shadow sausage*, first described by Cyrus Adler (1967), which is an interrupted shadow cast by a stick having a part submerged in water and another above the water surface. Here is a description of the phenomenon in a more recent paper:

> "A curved meniscus of the water surface is formed in the vicinity of its attachment to the branch. Sunlight casts a shadow of the branch on the sediments at the bottom of the stream. Because of refraction at the curved meniscus, the shadow con- sists of two disjoint segments corresponding to the portions of the branch above and below the water surface. The end of each shadow segment is pointed. The two shadow segments were

named the shadow-sausage effect by the author of Ref. 1, *since they resembled a pair of linked sausages*." (Lock *et al* 2003, *emphasis mine*)

As revealed by the emphasized passage, the reason this shadow consisting of two disjoint parts has been named a "shadow sausage" merely the *resemblance* to a connected sausage pair, so it is doubtful that we do actually perceive it as one shadow the way would perceive a sausage as one object. Hence, it looks as though amodal completion of shadow appearances is an exception rather than the rule.

VIII

Even though amodal completion of shadow appearances does not, as such, constitute a strong enough basis for the view that shadow perception is ultimately and basically no different from ordinary object perception, we do get a pretty strong indirect, empirically founded argument for the opposite view, which I have been propounding here, based on this phenomenon in the context of inquiring about the role shadow perception plays in the amodal completion of ordinary object appearances. We will see that the way shadow perception contributes to this is indicative of the likely truth of my view, namely, that seeing a shadow is a degree of failing to see a surface or a volume.

Tomonaga and Imura (2010) compared humans to chimpanzees in terms of the role played by shadows in their visual systems when it comes to amodal completion of appearances of shapes. The subjects were asked/trained to complete a search task in which they had to find the *pacman* figures (a disk with a missing quarter) from a set of distractors (disk and filled square shapes). Amodal completion occurs in both species when the pacman's appearance coincides with that of a disk partly occluded by a square. In these cases our visual system completes the pacman into a disk, hence, the search time for the pacman increases. As mentioned above, this phenomenon occurs in both humans and chimpanzees. However, when the pacman has an attached shadow along its square edge, such that the shadow appears as cast on the corresponding edge of the square

figure, humans and chimpanzees diverge in terms of search time needed to find the pacman:

> "All participants responded slowly under the occluded trials (without cast shadows) and rapidly under the control trials, replicating the previous results. The response times under the gap trials, however, varied between species and across conditions. The shadow cast by Pacman on the square prevented perceptual completion in humans but not in chimpanzees." (2010: 3)

The way our visual system, unlike that of the chimpanzee, treats shadows indicates that it is "interested" in what lies underneath the shadow. In treating shadows as shadows rather than as black objects, our visual system automatically aims at something beyond the shadow itself, and which is obfuscated by the shadow. For the chimpanzee's visual system, on the other hand:

> "(...) the black area functioned as the other "object." If the black "object" were superimposed on Pacman in addition to the square, the results would have been the same as those obtained when this object was not presented. However, if the object were superimposed on the square, this object would have been integrated with the square, not with Pacman's shadow, and perceptual completion would have been achieved." (2010: 5)

This adds empirical evidence to our view that in perception shadows do not function as ordinary perceptual objects, which our visual system aims at, but rather as "veils" between us and what our visual system aims at, namely, the surface or volume obfuscated by the shadow.

References

Adler, C. (1967) "Shadow-sausage effect," *American Journal of Physics* 35: 774 – 776.

Aranyosi, I. (2007) "Shadows of constitution," *The Monist* 90(3): 315 – 332.

Casati, R. And A. Varzi (1994) *Holes and Other Superficialities*, Cambridge, MA: MIT Press.

Dretske, F. (1969) *Seeing and Knowing*. Chicago: The University of Chicago Press.

Lock, J. A. (2003) "Analysis of the shadow-sausage effect caustic," *Applied Optics* 42: 418 – 428.

Sayan, E. (1996) "A mereological look at motion," *Philosophical Studies* 84 (1): 75 – 89.

Sorensen, R. (2008) *Seeing Dark Things. The Philosophy of Shadows*, Oxford: Oxford University Press.

Tomonaga, M. and T. Imura (2010) "Pacman in the sky with shadows: the effect of cast shadows on the perceptual completion of occluded figures by chimpanzees and humans," *Behavioral and Brain Functions* 6 (38): 1 – 9.

Westphal, J. (2011) "Silhouettes are shadows", *Acta Analytica* 26 (2):187 – 197.