

Seeing through eyes, mirrors, shadows and pictures

Helen Yetter-Chappell¹ 

Published online: 11 July 2017

© The Author(s) 2017. This article is an open access publication

Abstract I argue that we can see in a great many cases that run counter to common sense. We can literally see through mirrors, in just the same way that we (literally) see through our eyes. We can, likewise, literally see through photographs, shadows, and (some) paintings. Rather than starting with an analysis of seeing, I present a series of evolving thought experiments, arguing that in each case there is no relevant difference between it and the previous case regarding whether we see. In a sense, my arguments can be thought of as akin to the Extended Mind Hypothesis (Clark and Chalmers 1998). But instead of arguing that our minds can extend into the world, I argue that our sensory organs can extend into the world. Among the things that emerge from this discussion are (1) that—contrary to Currie (1995) and Carroll (1996)—seeing an object *O* doesn't require being able to locate *O* with respect to yourself, (2) that—contrary to Sorensen (2008)—we can see objects by seeing their shadows, and (3) that—contrary to Walton (1984)—it doesn't matter whether the causal relation between *O* and yourself is mediated by beliefs.

Keywords Seeing · Photography · Perception · Shadows

1 Introduction

When you gaze at the night sky, you see thousands of stars. When you go to the ballet, you see dancers gliding across the stage. In these cases, you see through (by way of) your eyes. Could you see the dancers through a mirror in the same literal way? Might you literally see the stars through an IMAX theatre, or through a star chart?

✉ Helen Yetter-Chappell
helen.yetter-chappell@york.ac.uk

¹ University of York, York, UK

In the spirit of Kendall Walton (1984, 2008), I'll argue that we can see in a great many cases that run counter to common sense. We can literally see through mirrors, in just the same way that we (literally) see through our eyes. We can, likewise, literally see through photographs, shadows, and (some) paintings. I'll present a series of evolving thought experiments, arguing that in each case there is no relevant difference between it and the previous case regarding whether we see. In a sense, my arguments can be thought of as akin to the Extended Mind Hypothesis (Clark and Chalmers 1998). But instead of arguing that our minds can extend into the world, I'll argue that our sensory organs can extend into the world.¹

Among the things that will emerge from this discussion are (1) that—contra Gregory Currie (1995) and Noël Carroll (1996)—seeing an object *O* doesn't require being able to locate *O* with respect to yourself,² (2) that—contra Roy Sorensen (2008)—we can see objects by seeing their shadows, and (3) that—contra Kendall Walton (1984)—it doesn't matter whether the causal relation between *O* and yourself is mediated by beliefs.

The structure of the paper is as follows. Section 1 makes the case for the ability to (literally) see through mirrors, photographs, and film. Section 2 considers four objections to my arguments: (i) that seeing an object *O* requires seeing *O*'s relation to oneself, which is violated by cases of photographs, films, and some mirror set-ups, (ii) that seeing “through” photographs and films requires too many awarenesses: an awareness of the photograph as well as the object, whereas paradigm instances of seeing merely require an awareness of the object seen, (iii) that my argument structure—comparing series of cases and arguing that there's no point that makes a difference between seeing and not seeing—is methodologically flawed, and (iv) that my arguments can extend to shadows as well as to mirrors and photographs, but we don't want to say that we see objects through their shadows. I reject the first three arguments and embrace the implication of the fourth (that we can see objects through their shadows). Section 3 further extends the cases in which we can see. I argue that it's irrelevant whether the relation between the agent and the object seen is mediated by beliefs. It follows from this, together with the arguments from Sect. 1, that we can literally see through some paintings (and other handmade pictures). Section 4 wraps up, offering a way of reconciling my conclusions with commonsense: While my arguments show that we *can* see through mirrors, shadows, and film, they don't show that we *must* see “through” these mediators. I consider some additional criteria that must be met to genuinely see.³

¹ One may plausibly be able to develop a similar account of other sensory modalities. While this paper will focus exclusively on sight, it's an interesting question how far the arguments developed here can be extended: Do we literally hear through telephones and radios? Can you hear *George Gershwin* playing his own music by listening to a recording? Do phones and mp3 players extend our sensory organs into the world and across time? I'll set these questions aside, though considerations of how the arguments extend to other sensory modalities may prove illuminating.

² But see footnote 10 for some complications concerning whether the location of the agent is taken to correspond to the location of her brain/body or (partially) her sensory organs.

³ This paper does not attempt to present a complete theory of object perception. Just as the best work in applied ethics develops important and surprising insights regarding how we should think about particular moral cases, without either presupposing or seeking to establish a complete moral theory, so the ambition

2 Seeing through

I see a hawk outside my window. In order for me to see the hawk, light waves must bounce off of the hawk, traveling towards my eyes. These light waves are focused by the lenses of my eyes, and directed back towards my retinae. My retinae are sensitive to light, rather like the film of a camera. The impact of the light hitting them triggers neural activities that process the information, ultimately giving me a conscious awareness of the hawk.

Suppose that my brain had a redundancy in it, such that the image projected on the retina was again projected onto a second retina before being processed. (We might imagine the first retina functioning something like a mirror, directing the light on to the second retina).⁴ This would be inefficient, but it would not make it the case that I failed to see the hawk. Now suppose that this process is moved outside of my head. One of the two reflections of the light waves takes place three feet in front of me (in a mirror), rather than inside of my head. It seems this should also make no difference to whether or not I see the hawk. Surely it's mere chauvinism to privilege things going on inside of our skulls. So we can literally see objects in mirrors. We do not merely see reflections that we interpret as saying something about the non-mirror-world. (This seems quite natural. Ballet dancers practice in rooms lined with mirrors so as to see their bodies without needing to distort their placement. Ballet teachers frequently make demands such as "Look at your hips!" by which they do not mean that one should tilt one's head down to assess one's placement).

Today, people applying makeup on a train are as likely to use a "mirror app" as an old fashioned mirror. Whereas a mirror would take light waves bouncing off of an object and reflect them back in the other direction, the mirror app requires some more complicated processing involving a camera and computer coding to preserve the information given by the incoming light. But the information is preserved just the same.⁵ The fact that the route to the "reflection" of the image is slightly less

Footnote 3 continued

of this paper is to develop important and surprising insights regarding particular cases of putative perception. In this respect, it could be thought of as a paper in the "applied philosophy of perception". Of course, while not itself in the business of developing a theory of object perception, this paper will—if its central claims are correct—have important implications for the correct account of seeing.

While these arguments should be of interest to anyone who cares about visual perception, they will be of special significance to philosophers who think that perception has distinctive epistemic importance. For instance, many philosophers think that we're non-inferentially justified in believing the deliverances of our senses. If, as I argue, we can literally see objects through shadows and hand-made pictures, this will extend the range of cases in which our beliefs can be non-inferentially justified. (Note: I simply flag this as one potential implication, given other commitments one might hold. It is not my task here to argue that perception has such distinctive epistemic significance).

⁴ Cats and other animals that are active at night do effectively have mirrors built into their eyes. The tapetum lucidum is a layer of cells that functions as a mirror, directing stray light at the eye's rods (and causing cats' eyes to seem to glow in the dark). But unlike our imagined case, the photons don't all hit the tapetum lucidum prior to hitting the retina.

⁵ One might insist that we get more information from the mirror, as binocular disparity gives us depth cues in this case that are absent from the video screen. But surely this information is not essential to

direct doesn't seem to change the phenomena or make it any less true that you *see* yourself when you use your mirror app to fix your eyeliner.⁶

This again seems quite natural. Just as those of us with older cars rely on mirrors in backing up to see if there's anything behind us, those with newer cars see behind themselves using screens and cameras to project the view from the rear of the car. We might even treat the mirror and the video screen in an identical fashion. Perhaps the ultra-high-definition, ultra thin video screens of the future will be completely indistinguishable from mirrors (unless you smash them open to inspect the inner workings). These video screens will preserve the same information about the back of our cars, and will be psychologically processed by our brains in the same way mirrors are today. Surely we must say the same things about whether they facilitate literal seeing.

Note that there's always a time lag between what we see and our experience of it. The most extreme example of this is in the case of stars. We (literally) see stars. When we do so, we are seeing objects that may no longer exist, and seeing them as they were in the past. But the same is true when I look at the hawk. It takes time for light to get from the hawk to my eyes, so that in some sense as I look at the hawk, I am seeing is the world as it *was*. It takes more time for light to get from the hawk to a mirror and then to me, and perhaps slightly more time for light to reach my phone, be processed, and be spit back at me. But these are differences in degree that should not be able to make the difference between seeing and not-seeing.⁷

Whether the time delay is intentional or not should also make no difference. A mirror app operating on a very old phone, such that there is a noticeable lag between your movement and the display on the screen, will facilitate seeing just as a mirror app on a new and speedy phone. There seems little basis for thinking that an app on

Footnote 5 continued

seeing. A one-eyed pirate can see his peg leg—either directly or through a mirror—though he lacks information from binocular disparity.

⁶ One might insist that there's an important difference between the mirror and the mirror app: In the old fashioned mirror, it's the very same photons that hit the object which later hit your retina (setting off the brain's processing of this information). In the case of the mirror app, while all the same information is preserved, it's different photons that ultimately hit your retina. Suppose the physics were such that this is true. While this would allow us to carve a line between cases where the photons were the same and those where they were different, it's strange to think that this line represents an important barrier. First note that if seeing required the very photons that hit the object to ultimately hit your retina, we wouldn't be able to see objects by seeing their silhouettes. But intuitively, e.g. when a hunter sees a silhouette of a moose, backlit on a mountaintop by the moon, he *sees* the moose. One might try to patch this view by arguing that *either* there must be an unbroken connection between the photons traveling from the object to the eye *or* we must detect the absence of photons from the object by virtue of an unbroken connection between photons passing immediately adjacent to the object to the eye. But this also can't be right. Gregory Currie imagines a screen that—while preventing photons from passing through it—registers the precise patterns of light that hit it and instantly emits qualitatively identical patterns of light from the other side. "Looking through the screen would be just like looking through a window at an object beyond. The screen interrupts the light ray, but I think we would say that when we look at it we see the object on the other side (assuming the mechanism preserves [relevant] counterfactual dependence)." (Currie, 60) There is nothing crucial about unbroken photon links.

⁷ Section 2.3 defends the claim that (in the cases described) differences in degree cannot yield a difference in kind.

a device with sufficient processing power, but that was deliberately engineered to include a lag, would differ regarding its facilitating seeing.

We've so far concluded (1) that seeing an object *O* in a mirror literally amounts to seeing *O*, (2) that there's nothing special in this regard about the particular causal mechanism at work in a mirror—a mirror app or a surveillance video would facilitate seeing to the same degree, and (3) that it's not essential to seeing that our visual experiences be simultaneous (or near simultaneous) with what we're seeing. It follows that films (regardless of when they were made) can facilitate seeing. When I watch a video of John F. Kennedy shaking hands with Nikita Khrushchev, I can literally see Kennedy and Khrushchev, just as a contemporary might on live television, through a mirror, or from an audience.

One might question whether this could be extended to photographs. As Gregory Currie (1995) has noted, seeing typically gives us temporally extended information that allows us to track objects over time. Photography, unlike film, doesn't share this feature. But while this it's true that seeing is typically temporally extended, this shouldn't be taken as a requirement for seeing. When you're in a room where strobe lights offer the only illumination, you see the world as though a series of still images.⁸ People are here, then they're shifted slightly, then their expressions are slightly altered. You know that people must be moving, but this is an intellectual knowledge, akin to the knowledge that there was motion occurring when we look at a series of photographs. While it might seem that you can still track objects over time in the strobe light room, this needn't be the case. Imagine that there are two identical balls in the strobe light room. It's clear that we see the balls, despite the fact that we cannot track whether they have changed places between pulses of light.

Alternatively, we might imagine a case where there's only a single momentary flash of light. Perhaps your power has gone out in a storm. It's pitch black. Then there's a single flash of lightning, just long enough to give you a split-second glimpse of the murderer in your living room. Clearly you see the murderer, despite not having any temporally extended information about him or ability to track him over time. So the temporally instantaneous nature of photographs doesn't pose a barrier to our ability to see through them. When I look at a photograph of Kennedy and Khrushchev shaking hands (perhaps extracted from the video discussed above), I can literally *see* Kennedy and Khrushchev.

The next thing to note is that (complete) accuracy is not essential for seeing. Vision often fails to reveal aspects of what we see. Normal humans don't see UV light; color blind people fail to see red and green as distinct colors; a person standing in a fog or at a great distance fails to see details. But my failure to see UV light doesn't mean that I fail to see the sun; the color blind person's inability to distinguish red from green doesn't mean that he fails to see the strawberries in the bowl; your inability to see the details of the deer's antlers on a foggy night don't mean that you fail to see the deer. In addition to our visual representations' failing to

⁸ The first time I experienced this, I thought that I was suffocating: My brain seemed to process the fact that the visual input it was receiving wasn't temporally extended as evidence that it couldn't engage in temporally extended processes such as breathing.

include certain aspects, they may positively misrepresent aspects. The (straight) stick in water looks bent to me; nevertheless, I still see the stick.

There is doubtless a great deal to be said about the extent to which seeing must reflect the world. But I don't think that anything that can be said here will bear on the relationship between "normal" seeing⁹ and "seeing through". Just as someone without color vision might nevertheless see John F. Kennedy, so we can see him through black and white photographs. Just as you might see a deer in the fog, so you might see your face steamed up bathroom mirror. Distorting mirrors and films in which all the actors are elongated to look taller and more beautiful, may likewise facilitate seeing.

3 Objections

3.1 Seeing and self-location

But one might think that all of the tiny steps that I've taken away from normal seeing, have resulted in a very big jump away from normal seeing, and hence away from anything recognizable as *seeing*. This jump first appears in the case of mirrors, but shows up clearly in the case of film and photography.

It's natural to think that if we can see an object "through" its reflection in a single mirror, adding a second mirror or a third mirror will not change the fact that we see the object. After all, we're just adding more of something that doesn't make the difference between seeing and not-seeing. But we can construct a collection of mirrors in such a way that we can draw a sharp line between normal seeing and the through-mirror seeing. Normally, when I see O, I not only see O's relation to other objects in the scene, I also see O's relation to *me*. This might also be true in the case of a single mirror. But given a sufficient number of mirrors, placed at the right angles, I might see a reflection in a mirror, but have no idea where the object causing that reflection is located relative to me. Likewise, in a photograph or film, I have no visual awareness of where the objects are relative to me. So there is a sharp line between complicated mirror cases (and photographs/film) and normal vision. The question is whether this line implies that we do not *see* in the former cases. Currie (1995) and Carroll (1996) think that it does: "I submit that we do not speak literally of seeing objects unless I can perspicuously relate myself spatially to them—i.e. unless I know (roughly) where they are in the space I inhabit." (Carroll, 62)

But this egocentric requirement can't be right—at least not on a strong interpretation that makes trouble for seeing through photographs and film. Imagine that your eyes have a horrible disease that is causing them to gradually dissolve. If nothing is done, you will not only go blind, you will risk your diseased eyes releasing a toxin into the rest of your body. Your eyes must be removed.

⁹ By "normal seeing" I mean the sort of seeing that occurs through human eyes, without the aid of mirrors, photographs, film, or other contraptions—what we typically think of as paradigm instances of seeing.

Fortunately, scientists have created prosthetic eyes that function identically to your current eyes. Unfortunately, they're a little too large to fit inside of your skull. They must be worn in a special device strapped to your head. This device will wirelessly transmit the information it receives to your brain in just the same manner that your eyes currently do. You have the operation, and (outfitted with prosthetic eyes) you go about your life just as before. The world appears unchanged to you.

In this case, you clearly continue to see the world. The material the eyes are made out of and the fact that they now relay information to the brain wirelessly, rather than via continuous physical channels, doesn't make it the case that you don't see.

Now suppose that the device that holds your prosthetic eyes is bulky and uncomfortable to sleep in. So you set the eyes by your bedside at night. A friend—thinking that it would be amusing—sneaks into your room and steals your eyes, depositing them on his kitchen table. When you wake up in the morning, you will not have gone blind. You might scream and wonder whose house you're seeing. But intuitively, you are seeing *someone's* house. Had you been awake at the time they were stolen, it would not have been the case that, as the thief ran away, there was some point at which you ceased to see.

You see despite the fact that there's a sense in which you can't locate objects relative to yourself—you don't know where objects are relative to your brain/body.¹⁰ There is, however, another sense in which you do can locate objects relative to yourself—you know where the objects you see are relative to your sensory organ, your eyes. But this can't draw a line between the remote prosthetic eyes and photographs, as something analogous is true in the case of a photograph. You may not know where the objects depicted in the photograph are relative to your brain/body, but you know where they are relative to the camera's lens.

But in the case of the photograph, there are two relations. There's the relation of the object to the camera lens (which we know) and there's the relation of the object to the eye (which we don't know). The objector might think that it's essential for seeing that we have knowledge of the second relationship. This would draw a line between normal seeing and "seeing through" photographs and complicated mirrors. But this can't be right either. We can build on the remote prosthetic eye case to imagine that the eye includes a double retina—light hits the first retina (which is somewhat redundant), and then is reflected onto the second retina, from which the information is given to the brain. The second retina is detachable, and is capable of relaying information on to the brain from wherever it happens to be. Your brain/body, second retina, and the remainder of your prosthetic eye might all wind up in different places (if you have enough prankster friends). Even so, it wouldn't be the case that you don't see things.

But this is precisely analogous to the camera lens and the eye. In both cases, we're visually presented with the relationship between O and the first detector of

¹⁰ I'll remain neutral as to where the agent is located. The agent might be taken to be located (i) with the seat of her consciousness, (ii) wherever she has a body that is capable of manipulating the world, or (iii) (at least partially) where her sensory organs are located. If the agent were located (partially) with her sensory organs, then my arguments would extend us *as agents* into the world in quite radical ways. This egocentric requirement is preserved, but not in a way that creates trouble for seeing through photographs.

light (the camera/first retina). And in neither case are we visually presented with the relationship between O and the second detector (the eye/second retina). So if we think that we would see things even were our eyes discovered to have second floating retina, we should think that our eyes can see things in the case of photographs and complex mirrors. Or, at least, we shouldn't think that our inability to see the relationship between O and our brains/bodies poses a challenge to our ability to see objects through photographs and complex mirror structures.¹¹

This is not to deny that there is a sense in which seeing is egocentric. Seeing always presents the world to us from a particular vantage point, rather than an impartial God-view. (This is true as much in the case of photographs and the prosthetic eye as in cases of normal vision). The upshot of the detachable prosthetic eye is rather that the vantage point needn't be the location of the agent (or, more precisely, it needn't be the location of the agent's seat of consciousness), but merely of the initial sensory input.

A second upshot of the prosthetic eye example is to illuminate some restrictions on the connection between seeing and action. Once your prosthetic eyes have been stolen, your seeing a tiger will not guide you in running away from the tiger. Your desire to see some writing more clearly will not guide you in moving closer to the text. These cases suggest that seeing some object O doesn't require your being able to act/move in relation to O.¹²

¹¹ Note that this argument applies equally to Jonathan Cohen and Aaron Meskin's non-doxastic account of seeing, according to which "x sees y through a visual process z only if z carries information about the egocentric location of y with respect to x" (Cohen and Meskin 2004, 201).

¹² Bence Nanay (2010) offers an argument in this spirit against the possibility of seeing-through photographs. Nanay argues that we cannot see through photographs, on grounds that seeing requires that there be "at least one way for me to move such that if I were to move in this way, my view of the perceived object would change continuously as I move" (2010, 468). We (apparently) can't continuously move in such a way that our view of the photograph's object continuously changes. So we can't see objects by way of photographs.

The extent to which my arguments are compatible with Nanay's requirement on seeing depends on various interpretational details. But however we interpret his requirement, I don't think that it poses a substantial challenge to seeing through photographs/film. (i) Nanay's requirement is implicitly committed to the idea that seeing must be temporally extended, which the strobe light case already gave us reason to reject. (ii) If motion is taken to be a deliberate action, distinct from *being moved*, the prosthetic eye example gives us reason to reject the requirement (as above). But even if the requirement is interpreted in a passive way as holding that there must be at least one way *for me to be moved* such that if I were moved in this way ..., the argument runs into difficulties, concerning what it takes for "*me to move*". (iiia) If, for an agent to move, the location of the seat of the agent's consciousness must move, then the prosthetic eye case is a straightforward counterexample. We can see an object O even when moving our seat of consciousness in relation to O doesn't change our perceptions. (In the prosthetic eye example, the motion of my brain in relation to O is irrelevant— it's the motion of the prosthetic eyes that's important.) (iiib) On the other hand, suppose that the movement of an agent simply requires movement of the part of the sensory organ at which there is initial sensory input. On this reading the prosthetic eyes are not a counterexample: There will plausibly always be ways that the eyes could be moved such that the agent's view of the perceived object would change continuously. But this way of understanding agent motion doesn't block the possibility of seeing through film. I've argued that video cameras function just as the prosthetic eye. There is no basis for treating the prosthetic "eye" as part of the agent's sensory system, but not the camera. And, just as with the prosthetic eye, there are ways that a movie camera could be moved such that the agent's view of the perceived object would change continuously.

But there are weaker senses in which our examples are compatible with a connection between seeing and action. It might be, for instance, that it's a precondition on our being able to make sense of spatial relations and three-dimensional objects that we have been able to move ourselves about in the world, and that this is crucial for seeing. (Susanna Schellenberg (2007) argues to this effect.) My prosthetic eye example does not show that an agent whose sensory organs had *always* been disconnected from their body (action-potential) would be able to see. It might be that someone whose prosthetic eyes were implanted and stolen at birth wouldn't be able to see, and that (correspondingly) someone who'd only ever received visual input by way of videos couldn't "see through" them. But this doesn't mean that *we* cannot see through them.¹³

Another possible restriction on seeing, which is compatible with the remote prosthetic eye examples, is that perceivers understand how their perceptions would change if their sense organs (were) moved. If someone were to pick up my prosthetic eyes and shift their relation to a teacup, I would not be surprised to find that the formerly circular-looking lip of a teacup now looked elliptical; nor would I be inclined to think that the shape of the cup had shifted. My reactions as I watch television are similar. When the camera moves in relation to a cup, I am not inclined to think that the cup has changed shape, but rather that the camera (like the prosthetic eye) has shifted position. So these weaker connections between perception and action are not threatened by our discussion.

3.2 Too many awarenesses

There's a second way an objector might try to distinguish "normal" seeing from cases involving mirrors and photographs. When I see my family through a photograph, I can be visually aware of two things: my family, and the photograph. By contrast, when I see my family face-to-face, I'm simply aware of my family. Perhaps this duality precludes seeing.

This insight is able to draw a line, separating off photograph, film, and mirror cases from "normal" seeing. But I don't think that this is a desirable line to draw. Imagine a creature who has introspective access to its visual cortex. It is able to

¹³ One might object: If we say that someone who'd never moved their sensory organs about the world at all would be unable to see, shouldn't we also say that someone who's never moved *this* sensory organ (a camera) about the world would be unable so see *using this sensory organ*?

I think it's clear that this would be the wrong thing to say. Insofar as it's plausible that having moved one's sensory organs about the world is a precondition on seeing, it is because this is plausibly necessary for being acquiring the ability to make sense of spatial relations, which in turn is plausibly necessary for seeing (Schellenberg 2007). To see that it's irrelevant whether *this* sensory organ (the camera) or some other one (your eye) was the one to have been moved through space, consider a case in which the person who receives the prosthetic eye has recently become completely paralyzed. She used to run, jump, play, and see perfectly normally through her eyes. She then becomes paralyzed and is fitted with a prosthetic eye which, alas, are fixed straight ahead. It seems clear that she continues to see through these prosthetic eyes, just as she would have had she simply become paralyzed and her original eyes become fixed rigidly ahead. It is irrelevant whether she ever moved *this* device for seeing about the world. What is relevant—if anything—is that she have been able to move *some* such device about the world, such that she can acquire the relevant spatial understanding.

consciously introspect and attend to the information that passes from its retina to its brain. Surely we don't want to say that this creature *can't see* simply because it has this additional introspective capacity. We want to say that not only can it see, it *also* has the capacity to reflect on the causal stream that makes vision possible. But this is just parallel to the situation we're in with mirrors and photographs, the difference simply lying in where along the causal chain we have access (and whether what we have access to is inside of our skulls).

One might try to argue that this is not analogous to our situation regarding mirrors and photographs, on grounds that when we see O through a mirror/photograph, we're not merely aware of two things—the image and the object O. Rather, our awareness of O is *dependent on* a conscious awareness of the image. We're consciously aware of the image and *thereby* aware of O.

But is it true that when I see my family in a photograph, my awareness of my family is *dependent on a conscious awareness* of the image? I look through my office window to the tree outside. I see the tree. While I could focus my attention differently—so that I am aware just of the window, or of both the window and the tree beyond it—I can also “look beyond” the window, so that the only thing I am consciously aware of is the tree. Likewise, when looking through a mirror: I can focus my attention so that I'm aware of the light reflecting off the mirror. But I can also attend solely to the scene reflected, so that this is all that I am consciously aware of. Thus, my awareness of the reflected scene doesn't seem dependent on my being consciously aware of the mirror's surface. The same is plausibly true of film: While I can attend to a movie so that I am consciously aware of the image on the screen as well as (say) Jack Nicholson's face, I can also attend simply to Nicholson, without conscious awareness of the light reflecting from the screen. Awareness of the object the camera is pointed at doesn't always seem to entail a dual awareness. If this is right, awareness of the object isn't *dependent on* conscious awareness of an image.

It is true that I am only aware of Nicholson's face because my brain is responsive to the light emanating from the screen. One might, in some sense, take this to involve an *awareness* of the light (even if it is not conceptualized as such). But this cannot help us to distinguish seeing through mirrors and photographs from uncontroversial cases of seeing. When I look at my cat, see him only because my brain is responsive to the image projected onto my retina. Thus in uncontroversial cases, as well as cases of seeing by way of films/photographs, I am aware of some object O only because my brain is responsive to some intermediary R (which might be loosely described as my being—nonconsciously, nonconceptually—aware of R). But this poses no barrier to seeing in uncontroversial cases. So it's difficult to see why more of the same should pose a barrier to seeing through films or photographs.

3.3 Differences in degree can yield differences in kind

The arguments I have offered take the following structure: (1) Present a series of cases, starting with uncontroversial cases of seeing. (2) Show that each new case merely differs from the previous in degree. (3) Argue that this difference in degree shouldn't yield a difference in kind: There's nowhere along the series of cases that

we can justify drawing a line between those that facilitate seeing and those that don't. (If a double retina doesn't preclude seeing, then a mirror shouldn't preclude seeing. If a small time lag doesn't preclude seeing, then a larger time lag shouldn't preclude it. If a simple causal mechanism can underlie seeing, a more complex one shouldn't preclude it).

One might object to the structure of these arguments, particularly to the move from (2) to (3). The principle "differences in degree can't yield differences in kind" is not generally a good principle. A single molecule of H_2O is not a liquid. But if you add another molecule, then another ... eventually you get something that is a liquid. These differences in degree add up to a difference in kind. When you take a person with a full head of hair, and you remove one hair, then another ... eventually you get someone who's bald. Again, differences in degree add up to a difference in kind. Why not think something similar could be true of the series of cases that I've offered? (Perhaps a simple causal mechanism can facilitate seeing, but if it gets too complex, it can no longer facilitate seeing. Perhaps a small time lag is acceptable, but once the time lag is too great—perhaps too great relative to the distance between the observer and the object—we can no longer count as seeing the object.)

In cases where differences in degree do yield differences in kind, one of three things seem to be going on, none of which captures the cases I've described. First, there are cases of vague predicates (e.g. baldness). In these cases, *every* difference in degree makes a small difference, and enough small differences add up to a big difference. Each step you take away from the clear-cut case of baldness makes them *less bald*. Every inch of height you add to a person's height makes them *more tall*. In these cases, there's a particular dimension along which we're evaluating, e.g. with the least tall person at one end, such that every step along this dimension (even from one paradigm tall person to another) makes one *more tall*. But in the cases involving seeing (e.g. when your eyes are replaced with prosthetic eyes, or your friend carries your prosthetic eyes further and further from your head), it doesn't seem that you are *less seeing*. So they can't be explained by analogy to vague predicates.

But not all cases involving small steps along a dimension are like this. Imagine a continuum of angles from (roughly) 1° to 179° , where each step along the continuum makes the angle *more acute*. While each step along the continuum makes a difference to *how* acute the angle is, it is not the case that these small steps add up to a difference in kind. Most of the steps along the continuum are irrelevant to *whether* an angle is acute. It is only the step from 90° to 89° that makes this difference.¹⁴

Might seeing be understood on this model: as standing along a continuum of subtly different cases, where only particular step(s) along the continuum make the difference between seeing and not-seeing? Perhaps. But even if so, this is not a threat to the arguments that I've presented. If seeing is like being an acute angle, there's some particular step we can take that transparently takes us away from seeing to not-seeing (much as the step from 89° to 90° transparently takes us from

¹⁴ Thanks to an anonymous referee for pointing out this complication.

acute to non-acute). But none of the steps that I've taken in the paper seemed to do this. It doesn't seem that *any* small difference in the duration of a time-lag makes a difference in kind. It doesn't seem that whether a mirror is inside of our outside of someone's skull makes a difference. While there certainly *could* be a step that transparently makes the difference between seeing and not-seeing, and draws a distinction between normal seeing and "seeing through" photographs, we haven't found it. And to search for such a difference is precisely the task this paper has pursued.

A third sort of case in which differences in degree yield a difference in kind involves emergent properties. As you add more and more H₂O molecules, the result doesn't gradually become *more liquid-like*. Nevertheless, these differences in degree add up to a difference in kind. In cases like this, it's transparent that when you arrange the low-level features of the world in the right way, the result is the emergence of a new macro-level phenomenon.¹⁵ This is precisely not the case with the examples I've given of seeing. For some number x , if we play out the behavior of x H₂O molecules, and then the behavior of $x + 1$ H₂O molecules, we can see that a new phenomenon has emerged. By contrast, when you play out a case with a time lag of $x + 1$ between an event and your visual awareness of it, and a case with a time lag of x , there doesn't seem to be any new phenomenon emerging. If we imagine a mirror that's inside of the head, reflecting light back onto the retina, and then imagine pulling that mirror outside of the skull, there doesn't seem to be any new phenomenon emerging.

So while it would be an over-generalization to say that differences in degree *never* yield differences in kind, this doesn't help to avoid the series of arguments I've given. The cases don't intuitively becoming *less seeing* with each step, nor do we have a new phenomenon that transparently emerges (or disappears) at any point along the cases.¹⁶

¹⁵ A single molecule of H₂O does not have the relevant functional profile to constitute a liquid, but once the relevant threshold is passed the molecules will constitute liquid water. To the best of my knowledge the precise number of molecules necessary for H₂O to function as a liquid at room temperature has not been pinpointed. But the minimum number of molecules necessary to constitute ice has: a minimum of 275 H₂O molecules are necessary to form the crystalline structure essential to water in its solid form. (Pradzynski et al. 2012).

¹⁶ Note that there may be cases that combine the different sorts of processes: Incremental changes can take us from a beautiful spring day, where Bob is sunbathing ... to Bob's untimely demise. Every fractional drop in temperature makes it *more cold* and (ultimately) *more deadly* outside. The incremental changes in temperature have incremental changes on the brain of Bob. At some point these incremental changes lead to a radical shift: Bob is dead. (Imagine that being alive is an emergent property, which Bob loses.) But this scenario simply involves a conjunction of the two described processes. We can break the scenario down into component pieces, each of which seems to involve either (i) gradual change, where every difference makes a difference, or (ii) a new structural property emerging/ceasing. This paper has proceeded by breaking down the differences between seeing objects directly and seeing objects by way of photographs along their dimensions of difference. In contrast to the scenario described above, none of these dimensions seemed to involve either gradual change in which every difference makes a difference, or the emergence of structural properties.

3.4 Shadows

A final challenge comes from shadows. It's intuitive to think that we see things "through" mirrors—it's perfectly natural for me to say that I see my hip placement ("in the glass") or that I see the child behind me (in my rear view mirror). By contrast, we don't think that when we see a shadow, we are thereby seeing the object that produced the shadow. Intuitively, when I see Bob's shadow protruding from behind a tree, I don't literally see Bob. I see Bob's shadow. Can we grant this intuition? Or must we conclude that, as with mirrors and photographs, we can see "through" shadows?

At first pass, it's difficult to see how we could be justified in treating shadows differently from reflections. Imagine that a shadow is projected onto the floor. The floor in a sense functions just as a mirror: A mirror is an intermediary in transferring information about the object. In the mirror, this information is given in the form of what light is reflected by the object. The floor likewise functions as an intermediary in transferring information about the object. It seems the only differences between the cases of shadows and reflections are (1) whether the information about the object is transferred by way of light reflectance or light blockage, and (2) the richness of the information about the object (which in the case of shadows reveals only outlines of opacity).

We know that the comparative lack of richness of the information shadows carry can't make the difference between seeing and not-seeing. I see in a fog despite the lack of detail; I see my desk when the room is dark, despite the fact that the colors are completely washed out by darkness and all I can pick out are outlines of comparative brightness.

One might think that there was an important distinction between "positive information" about what light is reflected by an object and "negative information" (information about absences, light blockage). If this were right, the difference between mirrors and shadows would come before the point at which light hits the mirror or the ground—it would be the sort of information that was initially transferred that distinguished them. But I don't think that we want to treat information transmitted by way of light reflection as different from that transmitted by way of light blockage.

First, privileging information transmitted by light reflection seems like quite an arbitrary way of carving a line between seeing and not-seeing. It's conceivable that there be creatures who get all of their visual information by way of absences. Suppose that these creatures live in a highly reflective environment and their eyes emit white light. Their brains are wired to detect color, but color is processed in terms of absences-of-reflection of wavelengths. When they detect blue objects, what they're picking up on is not the positive reflection of blue light, but the absence of other wavelengths of light. It seems strange to think that we see, but these creatures—who perhaps even have visual phenomenology very much like our own—don't.

Moreover, privileging positive light reflection doesn't carve a line that we want to carve even in our own visual experiences. When we see Alfred Hitchcock's silhouette at the beginning of Alfred Hitchcock Presents, we think that we *see*

Alfred Hitchcock. When a hunter sees the silhouette of a moose standing on a mountain, backlit by the moon, we think that the hunter *sees* the moose. But silhouettes, like shadows, relay information by way light blockage.

Such cases in which we see objects by virtue of our sensitivity to *absences of light* are more common than one might think, as illustrated by Roy Sorensen's (2008) example of para-reflections. Suppose that you are looking at a mirror image of a chessboard. You intuitively see the chessboard through the mirror: equally seeing the black pieces and the white pieces. But while you are seeing the white pieces by virtue of their *reflection*—the light is positively reflected off of them—you are seeing the black pieces by virtue of their *failure to reflect* light. As Sorensen (138) puts it “The visual system does not treat darkness as privation of light. [Rather it has a] tendency to treat darkness as if it were substantial (or at least as substantial as light).” Every time you look at something black through a mirror, you are seeing the object by virtue of detecting an absence of light. So whether information is transmitted by light reflectance or light blockage doesn't make a difference for seeing.

But while a distinction between “positive” and “negative” information may not be able to account for the intuition that we do not see objects by way of their shadows, Sorensen suggests an alternative way to uphold this intuition. We might hope that his work could offer a way of carving shadows off from reflections, silhouettes, and the like.

There are several interrelated arguments that might be found in Sorensen. The first argument starts from the idea that we can only see objects by seeing their parts. Sorensen (41) thinks that “[b]ecause an intrinsic change in the silhouette constitutes an intrinsic change in the object, the silhouette is part of the object”. So we can see objects by seeing silhouettes. By contrast, shadows are not parts of their objects—we can change a shadow without changing the object (say, by putting another object in between the object and the light source). So we cannot see objects by seeing shadows.

We can quickly set this argument aside. Everyone can agree that shadows are not parts of objects, but this is beside the point. If shadows facilitate seeing objects, they do so in the sense that our retinæ facilitate seeing objects. Our retinæ are not *parts* of the objects that we see, though they enable us to see parts and (so) to see objects. Likewise, Bob's shadow is not part of Bob, though (if the analogy to mirrors is right) the shadow may enable me to see Bob's head, shoulders, and legs, and so to see Bob.

A second, related idea found in Sorensen (39–42) is that we can only see objects via genuine causal processes, not mere “pseudoprocesses”. When one ball hits another, causing the second to roll away, this is a causal process. But the *shadow* of the one ball hitting the other, and the second ball rolling away is not. The shadows are epiphenomena: Each moment of their existence is caused *not* by earlier moments of their existence, but by the objects they're shadows of. Silhouettes, we might think, are different insofar as they are actually parts of their objects. They aren't an additional things caused by the objects, but are simply the objects, backlit. When we see the silhouette of one ball hitting another, and the second rolling off, we are seeing (via) a genuine causal process, and so (if we can only see objects via

genuine causal processes) we can see objects by seeing their silhouettes, but not by seeing their shadows.

But if this were right, it would show far more than that we don't see through shadows. It would show that we don't see anything. The mechanism that allows us to see in uncontroversial cases of seeing crucially makes use of our retinæ. But the light patterns hitting our retinæ—like the patterns of darkness hitting the floor in a shadow—are an epiphenomenon caused by changing conditions in the world.¹⁷

But there is perhaps a more sophisticated development of this argument, building on the idea that seeing essentially requires *tracking* a causal process. Suppose Mary stabs Bob, leaving a knife sticking out of his chest. This is a causal process. Were Bob and Mary standing behind a screen, with backlighting, their silhouettes would track this causal process in the sense that a modification to Bob's silhouette (as by coming to include a knife in the chest) entails a modification of Bob. By contrast, shadows don't track causal processes in this way. A modification in Bob's shadow could be caused by Mary's placing a knife in space with the appropriate relationship between the knife, Bob, and the light source.

Here one might object that we could create the same sort of deception with a silhouette (standing with a knife in the appropriate position behind Bob). The reply must be something like this: In the case of the shadow, our careful positioning of the knife has literally *changed the shadow*. By contrast, our positioning of the knife in the silhouette case doesn't really change the silhouette anymore than it really changes Bob. (Bob's silhouette is, after all, part of Bob in a way that his shadow is not.) Thus, the knife positioning merely creates the *illusion* that Bob has been stabbed in the chest. But illusions don't preclude seeing, even if they preclude seeing certain features of the scene. So I can see Bob by seeing the (illusory) silhouette. By contrast, when Mary holds a knife up in the shadow case, she literally changes the shadow, as opposed to creating an illusion of difference. So the shadow fails to track Bob. So I can't see Bob by seeing the shadow.

But while the shadow may be changed, it's not clear this can do the necessary work. Suppose you hold a straight stick above a pond, and then put the stick into the water at an angle (so that it appears bent). The image projected onto my retinæ from the stick differs between the two cases. In the second case, the light rays are bent as they pass from the water into the air, causing a change in what is projected onto the retinæ. What is projected onto the retinæ does not merely appear to change (as one might think the silhouette merely appears to change); it really does change (just as the shadow really does change). Nevertheless, we see the stick. If we can see the stick (and if it counts as an illusion), then we should likewise agree that (1) what we do when we look at shadows can constitute seeing and that (2) what we see when we look at the shadow of Mary, with the knife apparently embedded in Bob's chest, is an illusion.

¹⁷ The light hitting our retina is epiphenomenal in the sense that one light patch doesn't affect another light patch, they are merely the byproducts of causal processes in the world. It is, of course, true that the light hitting our retina affects our brain. But the same is ultimately true of shadows.

What we see through shadows can leave out a lot of detail. It leaves out color and depth. And because of the sparseness of what it conveys, it fails to represent differentiations between layered objects. But we've seen that detail is not essential for seeing. Our visual impressions when we see objects can be incredibly sparse. So perhaps where our reasoning about shadows has gone wrong has been in mistaking *sparse tracking* (that gives us a muddy picture) for a *lack of tracking*. Shadows don't represent or track distinctions between objects in the foreground and those behind them. But this doesn't mean that objects aren't represented and tracked. (Colorblind people don't track the difference between red and green, but they still see apple trees.) Likewise, what our retina tracks is light reflected off objects as it is when it hits the cornea. This can sometimes cause us to miss out on things (like the bentness of the stick), but it doesn't mean that we don't see. So just as we see through our eyes, mirrors, film, photographs, and silhouettes, we can see through shadows. When I see Bob's shadow poking out from behind the tree, I really do see Bob.¹⁸

4 Belief mediation

I want to conclude with one final, surprising way in which we can see things. It's common sense to think that we cannot literally see objects through paintings. No matter how realistic the painting is, how closely it happens to match reality, and how sincere the painter was in his belief that he was capturing reality, we are not literally seeing through to the objects depicted in the painting. Kendall Walton, who famously defends the thesis that mirrors, photographs, and film are transparent, denies the transparency of handmade pictures on the grounds that they do not exhibit a *natural* causal dependency on the objects depicted. But once we have accepted the possibility of seeing through photographs, I don't think that we can deny that we can also see through paintings (or other media for which an agent's beliefs about the world are a crucial part of the causal story). (Cf. Lopes 2004).¹⁹

¹⁸ Iterated pictures and pictures of shadows pose a further complication. Can we see object O by seeing a photograph of a photograph of O? Can we see person P by seeing a photograph of P's shadow? I think in both cases, the arguments of Sects. 1–2 entail that we can. If we can (literally) see a television set through the photograph of a living room, and we can (literally) see Barack Obama through the television set, surely we can see Obama through the photograph of the television broadcast. If we can (literally) see the sidewalk and the shadows on it through a photograph, and we can (literally) see Bob through his shadow, surely we can see Bob through the photograph of his shadow. This seems especially intuitively compelling when we consider cases involving mirrors. A dancer might use two mirrors, angled appropriately to see the back of her head, ensuring that her French twist is tidy. We are inclined to say that we see the back of the couple through the convex mirror in Jan van Eyck's *Arnolfini Portrait*—just as inclined as we are to say that we see their front. (But see Sect. 4 and footnote 27 for a caveat: The conclusion that we should draw from my arguments is that it's *possible* to see through the iterated picture, not that all cases of iterated pictures actually do facilitate seeing through.)

¹⁹ Lopes (2004) also rejects Walton's attempt to draw a line between photographs and hand-made pictures, though on different grounds. He argues that handmade pictures are not belief-dependent. To draw something recognizable as Piccadilly Circus, one needn't *believe* that one is drawing Piccadilly Circus. Likewise, Lopes thinks, to draw something recognizable as having feature F, one needn't *believe* that the thing has feature F. Further supporting this, Lopes argues that whereas belief possession requires having a *concept* of the objects/features involved, one can draw a picture without such conceptual

Walton argues against the transparency of handmade pictures as follows:

Photographs are counterfactually dependent on the scenes they portray: if the scene had been different the photograph would have been different. The same is often true of paintings, in particular when the artist painted from life aiming to portray accurately what he saw. But ... a painting from life depends counterfactually on the scene because the beliefs of the painter depend counterfactually on it. The counterfactual dependence of a photograph on the photographed scene, by contrast, is independent of the photographer's beliefs. It is because a difference in the scene would have affected the painter's beliefs about what is there, that it would have made the painting different. But a difference in what is in front of the camera would have made the photograph different even if it didn't affect the photographer's beliefs. The painter paints what he thinks he sees. The photographer captures with his camera whatever is in front of it, regardless of what he thinks is there. (Walton 2008, 127)

But it's difficult to see why *beliefs* are so special. What is it in virtue of which counterfactual dependencies making use of beliefs can't facilitate seeing, whereas those making use of other complex causal processes can?²⁰ Most philosophers think that human brains are simply very complicated information-processing systems—beliefs, simply functional states that store information and play appropriate causal roles within our brains. But while your brain may be more complicated than your iPhone, it's hard to see how this additional complexity of the causal tracking mechanism could make the difference between facilitating seeing and not. Further, even if you think that beliefs are more than information-storage operating in a

Footnote 19 continued

content. One needs to be able to discriminate and respond to the relevant features of the environment, but this can be done with non-conceptual content.

It's controversial whether there is such a thing as nonconceptual content (McDowell 1994). Lopes (2004, 185) writes that the content we experience of pictures' design features is "obviously nonconceptual", as we can e.g. discriminate between shapes without conceiving of them as having x or y sides. But this doesn't make the case that the artist fails to conceptualize F, as opposed to employing a recognitional concept of F.

Setting aside the question of whether there is nonconceptual content, we might wonder whether this will validate only a subset of handmade paintings as facilitating seeing-through. Even if an artist *needn't* conceptualize as they create, what happens if they do so conceptualize? Billy forms a concept of the color of a patch of sky, as he is painting ("that shade" or "the color of the lilacs at grandfather's house"). Can we see through the resulting picture? It seems to me that we can. It's hard to see why it should matter whether Billy forms a concept of the color of the sky as he paints, or whether he forms the belief that the sky is the color of the lilacs in his grandfather's yard. Lopes's defense doesn't seem to support this stronger conclusion.

²⁰ We might imagine a variation on Ned Block's China Brain thought experiment (Block 1978), in which the population of New Jersey simulates the workings of a digital camera. Although the resulting picture would be counterfactually dependent on the depicted scene in a way that crucially relied on belief, it's hard to see why this should render us unable to see through the resulting pictures. (While this case differs in the exact role belief plays in the counterfactual dependency, it does show that a general rule that seeing cannot be facilitated by a tracking mechanism relying on belief would be a bad one. We could likewise imagine the population of New Jersey simulating a prosthetic eye, feeding information directly into a person's brain. If they see with an ordinary prosthetic eye, they surely see here).

certain way in a functional system (perhaps they have intrinsic intentionality that the information stored in an iPhone lacks) it is hard to see how this addition could disrupt seeing. Beliefs *do* carry information as part of functional systems, whatever else might be true of them.²¹

To flesh this out, let's construct a conscious agent, Harold, whose produces pictures mediated by his beliefs, but who is in all other respects exactly like a digital camera. Harold has no imaginative capacity beyond a fantastic visual memory. After eating lunch in a blue room, with a red table, and yellow flowers in a silver vase, Harold can remember the room in perfect detail. But Harold lacks any sort of recombination or distortion capacity. He can remember the room *exactly* as it was, but he cannot imaginatively switch the colors of the walls and the table, imaginatively elongate the table, or imaginatively leave out the stain on his dining companion's shirt. If you suggest that he says something other than the truth—say, that the flowers are pink—he will become terribly flustered, “But the flowers are yellow!” Now suppose that Harold is an incredibly talented painter, able to paint works that are qualitatively indistinguishable from photographs. Harold's imaginative deficit means that he is unable to deliberately paint anything that deviates from how he believes reality to be.²² Just as a camera is designed to perfectly reflect what it takes in from the world, so too is Harold's mind. The difference is that the images Harold produces are in part the result of his *beliefs*. When Harold sits before his canvas, looking out at a scene, he selects a particular shade of blue because he believes it to match the ocean. When he draws the curve of a woman's shoulder *just so*, it's because he believes that this is what most perfectly reflects the curvature of her shoulder. And, thanks to his amazing memory and powers of discernment, he is always right about these things.²³

²¹ Similarly, when I see my alarm clock through a mirror, what's relevant is that my experiences track the clock. If additionally it turned out that God intended for the light that bounced off the mirror to preserve the structure of the light that originally hit the mirror, this would not show that I do not see things in mirrors. Adding nonphysical intentions or beliefs on top of the causal processes and information tracking doesn't seem like it changes things regarding whether or not we see.

²² Though just as a camera could malfunction—causing colors to become switched around, or even imposing features that were not present in reality—so too could Harold's brain malfunction. What's relevant is that Harold cannot *deliberately* paint anything that deviates from reality. He is effectively a belief-possessing camera, whose beliefs function as intermediaries between his eyes and the paintings he creates. Just as the possibility of my hallucinating doesn't preclude my seeing (in good cases), and the possibility of a camera malfunctioning doesn't preclude my seeing through photographs (in good cases), so too the possibility of Harold having brain glitches or hallucinations doesn't preclude our seeing through his paintings (in good cases).

²³ Note that this is a visual counterpart to David Sanford's case of Chris the Amazing Human Hearing Aid, who listens to sounds and repeats them so that what you hear is qualitatively indistinguishable from the original sounds (as cited in Noë (2003)). Sanford argues that one does *not* hear the lecture by way of hearing Chris, whereas one does by way of a conventional hearing aid. Alva Noë (2003, 98) explains this by appeal to differences in the egocentric information presented by Chris and the hearing aid: “Normally ... your auditory experiences of the people across the table will change if ... you turn away, or you get up and walk across the room. ... [To] hear how [things] sound from [Chris's] vantage point is not to hear how they sound (or how they would sound) from your own vantage point.”

In light of Sect. 2.1, I think that we should interpret Chris's case as one in which we hear—from *Chris's perspective*—analogously to how we see when our prosthetic eyes are stolen. Likewise for Harold. What's distinctive in the cases of Chris and Harold is that the sensory prosthetic is another agent.

Suppose that we're confronted with a photograph and a (qualitatively identical) Harold-painting of Barack and Michelle Obama holding hands on election night, 2012. Both pictures are true to life—there really was a moment that, from the right angle, a human agent would perceive as qualitatively identical to the two images. Both pictures would have been appropriately different had the scene at that moment been different. For both pictures, we can infer from the fact that they exist that a scene qualitatively like them really did occur. (We know from the way that cameras are designed that they will only show some feature F if they were confronted with F. And we know from the way that Harold is designed that he will only paint F if he is confronted with F). How could it possibly make any difference to seeing whether the causal mechanism involved in tracking F involves the particular complex functional state known as belief, or merely a less complicated functional state of the sort used by the digital camera?

Walton writes that photographs differ from paintings in that, in the case of photographs “a difference in what is in front of the camera would have made the photograph different even if it didn't affect the photographer's beliefs”, whereas with paintings there will be a difference in the painting if and only if there is a difference in the painter's beliefs (Walton 2008, 127). This is true, but it's hard to see how it makes a principled distinction between photographs and handmade pictures. When he paints, Harold will produce an image of a unicorn if and only if his brain is in a certain functional state—one that could be described as believing that there's a horse with a single horn on its forehead in front of him. Likewise, a camera will produce an image of a unicorn if and only if it is in a certain functional state—a functional state that stores information as of a unicorn in view of its lens. Walton's statement of the belief-independence of photography has a direct parallel in the case of Harold's paintings. Suppose that, prior to the development of cameras, journalists had dragged Harold about and commanded him to paint. It would have been true to say that “a difference in what is in front of Harold would have made the painting different even if it didn't affect the journalist's beliefs”.


So there seems to be no principled basis for denying that we can see through images where the causal process producing the image involves belief. Still, Harold is a very peculiar painter. How far can we extend these lessons? Are there actual handmade pictures that we can see through?

First consider Anne. Anne is just like Harold insofar as she has a fantastic visual memory, with no further imaginative capacities and no ability to produce paintings that (intentionally) diverge from her beliefs about reality. Unlike Harold, Anne's artistic abilities are limited. She draws stick figure people, trees as rectangles

Footnote 23 continued

While it might seem odd to think that one agent could function as an extended sensory organ for another, it is difficult to flesh out a coherent basis for rejecting such possibilities. This could again be seen as paralleling the extended mind hypothesis, about which Clark and Chalmers (1998) write: “Could my mental states be partly constituted by the states of other thinkers? We see no reason why not, in principle. ... What is central is a high degree of trust, reliance, and accessibility. ... For example, the waiter at my favorite restaurant might act as a repository of my beliefs about my favorite meals (this might even be construed as a case of extended desire).”

(trunks) with circles on top (leaves), and shallow, curvy ‘V’s for birds. Anne was sitting with Harold and the photographer on election night and drew the same scene of Barack and Michelle holding hands.



When we look at Anne’s picture, can we see Barack and Michelle Obama in the same literal way that we can see them through Harold’s picture, the photograph, the live election night television coverage, the mirror, and our eyes?

Again, the conclusion that we literally see Barack and Michelle is surprisingly difficult to avoid. We have seen that detail and accuracy are not essential for seeing. We see the sun, though we don’t see UV light; we see the deer in the fog, though we don’t see the details of its face or antlers; we see our bedrooms in the dark, though we see no colors and only register comparative brightness; and we see the pencil in the water though we don’t see the straightness of the pencil. Anne’s painting may leave out a lot of detail, but there are some things that it genuinely and successfully does track. Given the reliability of the tracking mechanism and our arguments thus far, it’s hard to see how we could deny the possibility of seeing those accurately depicted aspects of the Obamas.²⁴

Both Harold and Anne are highly unusual, in that their lack of imaginative powers has left them unable to create paintings that deliberately vary from the

²⁴ More radically still, we might imagine an even more abstract drawing of the Obamas, in which the illustration simply consists of two vertical rectangles. Can we see (say) the upright posture of the Obamas through such a drawing? Or is it possible for too little information to be conveyed to facilitate seeing? I can’t see any basis for thinking that it’s impossible to see (very limited) things through such an image. But see Sect. 4 for a significant caveat: The mere fact that it’s possible to see the upright posture of the Obamas through such an image does not mean that we always *do* see through such images. It may be that (due to psychological quirks) we’re less inclined to process information from very abstract images as revealing the world to us in the way required for seeing.

world. Because of this, there's a sense in which they are both (like cameras) highly reliable mechanisms for channeling the world. In this way, they differ from most painters, who can (and do) construct and embellish. But this does not mean that we can't see through paintings created by those with imaginative powers beyond Harold and Anne's.

When I look at a photograph, I have a good reason to believe that the picture was created by way of a certain sort of process that preserved information in much the way a mirror or a double-retina might. Of course, I may not always be able to tell this. Imagine that white bubbles periodically appear in the air. I have a camera that—though usually highly reliable—occasionally leaves white circles in images, qualitatively identical to the bubbles. I might sometimes be unable to distinguish whether a picture taken with this camera is reliable in the relevant sort of way. Hence, I may sometimes be unable to tell whether I am seeing through the photograph to white bubbles in the air. But this epistemic problem doesn't mean that I'm not (in the good cases) really seeing, any more than the fact that I might sometimes have hallucinations that are indistinguishable from perception means that I don't sometimes really see. Likewise, the possibility that a painter have hallucinations or illusions that affect their beliefs and thereby affect their paintings does not pose a challenge to our seeing through paintings in the good cases.²⁵

Though most paintings might be made by artists who embellish, are unreliable, or simply don't aim to reproduce the world-as-it-is, for those works that are created in a manner akin to Harold's paintings, we can see through them. Intrinsic features of the paintings will not reveal to us which paintings facilitate seeing through. So we will sometimes be left unsure as to whether we are literally seeing the Obamas or not. But this no more poses a problem for seeing "through" paintings than hallucination poses a problem for seeing in normal cases.

Having granted that the causal mechanisms facilitating seeing can include beliefs, we find that other ways of seeing open up. It's clear that I can see by way of a prosthetic eye (even if the exact mechanism by which the eye preserves information differs from that of a normal human eye). If beliefs can be part of the mechanism through which information is transferred, then we can even construct cases in which a scientist directly stimulating your brain facilitates seeing. Your eyes have been hopelessly damaged in an accident. You sit in a small room with a doctor (who is determined to enable you to see) and a monkey. The doctor understands exactly how light is processed between the eyes and the visual cortex, and has become skilled in replicating this process. As the monkey moves about the room, the doctor stimulates the electrodes throughout your visual cortex in just the way she believes will transfer the information about the monkey's motion. The end result of this process is that you have visual impressions exactly like those the doctor herself is experiencing. Despite the fact that the doctor's beliefs are crucial in the transfer of visual information about the monkey, the doctor enables you to see

²⁵ Even cameras might be subject to something akin to hallucinations and illusions. A fancy digital camera might occasionally get its signals crossed and combine information from pictures it previously took, much as a painter might occasionally hallucinate. But this does not show that we fail to see through photographs in the good cases.

the monkey. She is effectively functioning as part of a large, diffuse prosthetic eye.²⁶ Even conscious belief-possessing agents can function as parts of our extended sensory organs.

5 Failing to see through

I've argued that no principled line can be drawn to separate normal seeing from certain information-tracking visual presentations making use of mirrors, pictures, and shadows. Despite my arguments, the reader may feel that it's quite clear that what we do when we look at photographs (or watch television or watch shadows cast on the ground) is very different from what we do when we look at our cat or a tree or a cloud. In fact, I agree, and I do not think that this is at odds with anything argued in this paper.

I have argued simply that we *can* see through pictures, shadows, and mirrors, not that we must (or even that we commonly do). I'll conclude by proposing a way to reconcile these conclusions with our ordinary intuitions, suggesting that it's not sufficient for seeing that one processes information tracking the world, leading to visual phenomenology. The information must be processed in such a way that we seem to be confronted with the (tracked objects in the) world. While there's no principled barrier to our processing information from photographs, shadows, and pictures in this way, contingent features of our psychology may lead us to process them differently.

To demonstrate the importance of how we process information from our environment, we might develop Michael Huemer's (2001) useful distinction between objects of awareness and vehicles of awareness. Suppose I'm chopping wood with an axe. The wood is the *object* of my action; the axe is the *vehicle*. Likewise, in perception, Huemer thinks, we can draw a distinction between the *object* of awareness (the thing perceived) and the *vehicle* of awareness (the means by which the object is perceived: e.g. a mental representation, or—perhaps—a reflection, photograph, or shadow).

This analogy is helpful, but I think it should be taken further. When I'm chopping wood with an axe, the wood and the axe are important. But there's a third thing that's important: what I'm doing with the axe. This third thing is so important that we can't make sense of what the *object* is without understanding what activity is being performed.

²⁶ This is contrary to Walton (1984, 2008), who contrasts prosthetic eyes to a case involving a neurosurgeon directly stimulating your brain in the way described. Walton thinks that in the case of the neurosurgeon, the person would not see, whereas if the same stimulations were being made to the agent's brain by way of a non-belief-possessing prosthetic eye, you would see. But again, it's hard to see why e.g. making the eye conscious or making it have belief-constituting functional states would mean that the person using the eye would fail to see. Currie (1995, 61–62) makes this point another way. It could turn out, as per Malbranche, that God is the essential mediator of all impacts the world has on us. I see the blue of the sky only because God benevolently ensures that my sensory experiences match my surroundings. If this metaphysical hypothesis were right, we would still see the world around us, even though our perceptions were all mediated through God's beliefs. The doctor in this case is functioning analogously to God.

Consider two things I might do with an axe and some wood: I might use the axe to cut the wood. This makes the wood the object. (The object of what? Of cutting.) Alternatively, I might use the axe as part of a ritual blessing of the Wood Spirit. This makes the Wood Spirit the object. (The object of what? Of blessing.)

Likewise, my brain might do different things with the information hitting my optic nerve. Consider two cases. In the first, I look at a sunset and *see the sunset*. In the second, my gaze is directed at the sunset, but I'm completely inwardly focused: *introspecting on my phenomenal experience* to the exclusion of seeing the sunset. In the first case, my brain processes information V from my optic nerve as revealing something about tracked feature of the world, W. This makes W the object. (Object of what? Of perception.) In the second case, my brain processes information V as revealing something about *me*, not about the world. This makes V the object. (Object of what? Of introspection.) As with the wood and the axe, the details of the action that is performed are essential to making sense of what plays the role of the object.

Now consider a case in which I'm looking at a mirror, where a floral arrangement is reflected. Intuitively, I might either use the mirror to look at the flowers, or I might attend to the reflection itself. Following the above pattern, it seems that my brain might process information M—coming to my optic nerve via the mirror—as revealing something about the flowers. This makes the flowers the object. (Object of what? Of perception.) Alternatively, my brain might process information M as revealing something about the mirror, not the flowers. This makes the reflection in the mirror the object. (Object of what? Of perception.)

The same might be said for the different ways that we could engage with a photograph, shadow, or hand-made picture. More is relevant to whether we see through a particular photograph than just the photograph and the tracking relations it (and the observer) stands into the world. It's also important what the observer *does* with the photograph: *how* they use the information transmitted to them by way of the photograph.

I've argued that there's no principled barrier to our seeing through mirrors, photographs, shadows, and pictures. We *can* see through these mediators, but it doesn't follow from this that we must or that we typically do.

A full account of how information is processed in perception is beyond the scope of this paper. But I want to conclude by gesturing at several reasons for thinking that contingent features of the way we process information from our optic nerve can be relevant to what, or even whether, we see.

Consider two children arguing. Clara says "I'm mad because you ate the last cookie." Fritz replies "I'm mad because you're mad because I ate the last cookie." Clara: "And I'm mad because you're mad because I'm mad because you ate the last cookie." While there's no principled barrier to going on with this indefinitely, contingent features of human psychology are such that we actually can't go much beyond three iterations.²⁷

²⁷ And insofar as intuitions may balk at many-layered iterated pictures (seeing through a photograph of a reflection of a shadow), this very psychological phenomenon may account for the reason.

While the details of this case are quite different from those we've been considering, this example effectively illustrates that there's more to the behavior that we engage in than considerations of metaphysical possibility. There's no principled barrier to carrying out a thousand layers of this iterated argument. We'd be mistaken if we tried to argue that this would be impossible. But that's not to say that we (even the most dogged children among us!) actually do engage in thousand-layered arguments.

Likewise, I've argued that there's no principled barrier to seeing through mirrors, photographs, (some) hand-made pictures, and shadows. But that's not to say that we always, or even frequently, do so. For there may be contingent features of our psychology that mean that we often *don't* process information from photographs (etc.) in the way necessary to see through them.

We shouldn't be surprised to find such quirks of our psychology affecting the way that we engage with certain mediating images—and hence what our brains *do* with the information they take in via these images—as there is evidence that we tend to think about different mediating images differently, in a way that correlates with how intuitive we find seeing-through the image.

- (i) Our intuitions balk at the idea of literally seeing through pictures and shadows, though we are perfectly happy with the idea of seeing face-to-face, and literally seeing through mirrors seems plausible to many people. This difference seems to go hand-in-hand with a difference between how readily we think of the mediating images (projections on retina, reflections, shadows, pictures) as things in their own right. While there isn't any important *metaphysical* difference between the images projected on our retina and shadows, this difference in how we *think of* them may signal a difference in the way our brains tend to process the information, and so in how likely we are to actually use them to see.
- (ii) We also seem less inclined to psychologically process information from still images as revealing the world to us. Again, this seems to be due to contingent features how our brains typically process visual information, rather than to any principled difference between momentary and extended visual inputs. We can see this by imagining cases in which it's plausible that we *would* see through a still image.

Imagine a room with an “Image Window” along one wall—a giant screen that shows precisely what's behind the screen, in such detail as to appear to be a window. (Perhaps cameras are set up directly behind the screen, ensuring that it functions just as a window.) Even independently of the arguments I've put forward in this paper, it's plausible to think that we see through this “Image Window”. Now imagine that this room is pitch black, except for a pulsing strobe light, resulting in visual experiences as though of still images. If we saw through the Image Window originally, we surely continue to see through it, just as we continue to see the decor of the room. But now imagine that the Image Window goes black in precise time with the strobe light, so that only still images are shown on this “Still Image Window”. The experience of being in the strobe light room with the Image Window

is *completely indistinguishable* from the experience of being in the strobe light room with the Still Image Window. In both cases, I think it's plausible that we would process our visual experiences in a way that facilitates seeing.

What lesson should we take from this? One possibility is that our brains implicitly assume that visual experiences presenting the world to us will be temporally extended. In order for us to see in cases where our visual inputs are not temporally extended, there must be some explanation for why this default assumption is overridden. (In the case of the Still Image Window, the strobe light might be thought to provide this explanation.) Regardless of the psychological explanation, what's important is that such a psychological restriction doesn't reveal anything about the nature of seeing, or about whether it's possible to see through still images. It simply reveals a contingent feature of our psychology, that makes us unlikely (in normal circumstances) to see through photographs and paintings. This does not speak to the more general question of whether it's possible to see through such images.

So there are differences in how we think about and engage with mediating images in cases that we intuitively think facilitate seeing (e.g. retinal projections) and those that we don't (e.g. photographs). I've suggested that seeing an object O requires more than merely standing in a certain tracking relation to O: we must additionally process the information we take in from O in the right way. We now have a basis for thinking that we often process information from different mediating images differently.

Importantly, the relevant information-processing seems to be subpersonal, rather than mediated by our explicit beliefs. When Gullible Gary goes to the zoo, and his friend tells him that all the animals are hallucinations, he comes to believe that the tiger is really a hallucination.²⁸ The fact that Gary *believes* that his visual phenomenology is a hallucination clearly doesn't preclude him from seeing. Fortunately, this is not in conflict with the basic idea that *what our brains do with the information they take in from the environment* (how they process the information) *matters for what and whether we see*. There is obviously some sense in which Gary's brain treats the information it takes in from the tiger differently, compared to a person who believes that he is seeing a tiger. But we needn't commit ourselves to thinking that *this* difference in processing is the relevant one to seeing. (Note that while Gary may *believe* that he's not seeing a tiger, there is a sense in which it may well still *seem to him as though* he is looking at a tiger ... in a way that—perhaps—it doesn't *seem as though* he's looking at a tiger when he looks at a photograph.)²⁹

We can see through mirrors, shadows, film, photographs, and (some) handmade pictures. But this does not mean that we always, or even typically, do so, as

²⁸ Thanks to an anonymous referee for raising this complication.

²⁹ Sorensen (2008, 185) illustrates a case in which it's natural to mistake an object and its shadow. We might likewise ask whether the confused viewer in this case fails to see the object by virtue of their confusion. It could be that this would follow from a complete theory of how our brains must process information to see. But—as with Gary—this isn't entailed by the core idea that what our brains do with the information they take in from the environment matters for what/whether we see.

contingent quirks of our psychology may mean that we sometimes fail to engage with these mediating images in the way required for seeing. Does this undercut the interestingness of the thesis that we can see through photographs, etc.? Not so. First, it's a surprising result that these mediating images *can be* (and, plausibly, sometimes are) seen through. Second, even if it were to turn out that we rarely (or never) engaged with these mediating images in the way required for seeing, we have a surprising analysis of *why* we fail to see in these cases, which is at odds with the conventional wisdom. There is no principled barrier to seeing through photographs (contra Currie (1995) and Carroll (1996)), shadows (contra Sorensen (2008)), or appropriately made handmade pictures (contra Walton (1984)). If anything prevents us from seeing through these mediators, it is a contingent quirk of how we sometimes engage with them.

Acknowledgements Many thanks to Richard Yetter Chappell, Sinan Dogramaci, Daniel Greco, Robert Hopkins, Boris Kment, Karen Lewis, Jim Pryor, Susanna Schellenberg, Kendall Walton, and an anonymous referee for *Philosophical Studies* for their very helpful discussion and comments on earlier drafts of this paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Block, N. (1978). Troubles with functionalism. *Minnesota Studies in the Philosophy of Science*, 9, 261–325.
- Carroll, N. (1996). *Theorizing the moving image*. Cambridge: Cambridge University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 7–19.
- Cohen, J., & Meskin, A. (2004). On the epistemic value of photographs. *Journal of Aesthetics and Art Criticism*, 62, 197–210.
- Currie, G. (1995). *Image and mind: Film, philosophy, and cognitive science*. Cambridge: Cambridge University Press.
- Huemer, M. (2001). *Skepticism and the veil of perception*. New York: Rowman and Littlefield.
- Lopes, D. (2004). *Understanding pictures*. Oxford: Oxford University Press.
- McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.
- Nanay, B. (2010). Transparency and sensorimotor contingencies: Do we see through photographs? *Pacific Philosophical Quarterly*, 91, 463–480.
- Noë, A. (2003). Causation and perception: The puzzle unravelled. *Analysis*, 63, 93–100.
- Pradzynski, C. C., Forck R. M., Zeuch, T., Slavíček, P., Buck, U. (2012). A fully size-resolved perspective on the crystallization of water clusters. *Science*, 21, 1529–1532.
- Schellenberg, S. (2007). Action and self-location in perception. *Mind*, 116, 603–632.
- Sorensen, R. (2008). *Seeing dark things: The philosophy of shadows*. Oxford: Oxford University Press.
- Walton, K. (1984). Transparent pictures: On the nature of photographic realism. *Nous*, 18, 67–72.
- Walton, K. (2008). *On pictures and photographs: Objections answered. Marvelous images: On values and the arts*. Oxford: Oxford University Press.