

On the Epistemic Value of Photographs

The photograph is the only picture that can truly convey information, even if it is technically faulty and the object can barely be identified. A painting of a murder is of no interest whatever; but a photograph of a murder fascinates everyone.

Gerhard Richter¹

Many have held that photographs give us a firmer epistemic connection to the world than do other depictive representations.² To take just one example, Bazin famously claimed that “[t]he objective nature of photography confers on it a quality of credibility absent from all other picture-making.”³ Unfortunately, while the intuition in question is widely shared, it has remained poorly understood. In this paper we propose to explain the special epistemic status of photographs. We take as our starting place (in Section I) Kendall Walton’s startling proposal that photographs are special because they are “transparent”⁴—that is, that they are special because, unlike other depictive representations, they enable us literally to see their depicta.⁵ Walton’s proposal has not convinced many; however, it has proven surprisingly difficult to say just what is wrong about the transparency thesis. In Sections II through IV we shall rise to this challenge and show why photographs are not transparent in Walton’s sense. Finally, in Sections V through VII we shall propose and defend a novel diagnosis of what is epistemically special about photographs.

I. TRANSPARENCY AND PHOTOGRAPHS

In saying that photographs are transparent, Walton means that visually attending to a photograph enables us to see something numerically

distinct from that photograph—namely, its depictum.⁶ For Walton, photographs are of a kind with mirrors, telescopes, and microscopes: they are prosthetic devices that enable us to see things that we could not see without them.⁷ Whereas these other prostheses help us to see things around corners, very distant things, and very small things, photographs enable us to see things that are spatiotemporally remote.⁸ Walton emphasizes that he means this proposal quite literally:

I must warn against watering down this suggestion, against taking it to be a colorful, or exaggerated, or not quite literal way of making a relatively mundane point. I am not saying that the person looking at the dusty photographs has the *impression* of seeing his ancestors—in fact, he doesn’t have the impression of seeing them “in the flesh,” with the unaided eye. I am not saying that photography *supplements* vision by helping us to discover things we can’t discover by seeing. . . . Nor is my point that what we see—photographs—are *duplicates* or *doubles* or *reproductions* of objects, or *substitutes* or *surrogates* for them. My claim is that we *see*, quite literally, our dead relatives themselves when we look at photographs of them.⁹

Why does Walton insist that photographs are transparent? He believes that there are significant similarities between the way that photographs provide visual experiences and the way that ordinary vision provides visual experiences. For one, photographic images are counterfactually dependent on the scenes they represent; for example, had your ancestor been smiling rather than frowning, the photograph of her would have looked different. For another, and unlike realistic paintings and drawings (where such counterfactual dependency may hold), this

counterfactual dependence is not mediated by the intentional states of any intermediary agents. As Gregory Currie puts it, there is a “natural dependence” of photographs on the scenes that they depict.¹⁰ Finally, photographs also preserve real similarity relations between objects: like ordinary perception, confusions about photographic representations (i.e., with respect to what they depict) tend to be linked to real similarities between objects.

For Walton, then, photographs are transparent but paintings are not. Moreover, he argues, this difference makes an epistemic difference—for example, it explains why the appearance of photographs, but not that of paintings, supports counterfactuals about the appearance of the depictum. In addition, it explains why we often treat photographs as evidence (both formal and informal), whereas we are resistant to treating paintings and drawings as such.

We believe that Walton’s proposal does highlight certain important features of photographs that are worth capturing. However, it has the significant defect that its core thesis—that of the transparency of photographs—is (to put it gently) highly counterintuitive.¹¹ But just what is wrong with this thesis? In particular, if we are to deny the thesis, we owe an explanation of what it is about photographs that makes them nontransparent, given that there are other visual prostheses, such as mirrors, microscopes, and telescopes, that are transparent. This, then, is Walton’s challenge to those who reject the transparency thesis: explain the relevant difference between photographs, on the one hand, and mirrors (etc.), on the other.

II. EGOCENTRIC SPATIAL BELIEFS

To motivate our own answer to Walton’s challenge, it will be useful to begin with a proposal that has been suggested by a number of authors,¹² and that turns on an appeal to visually represented spatial information. The idea here is that a necessary requirement for *x*’s seeing *y* is that *x* represents information about the spatial relations between *x* and *y*.¹³ This requirement, it has been suggested, effectively draws a line in the sand between uncontroversial examples of transparent visual prostheses on the one hand, and photographs on the other:

With ordinary seeing, we get information about the spatial and temporal relations between the object seen and ourselves. . . . Photographs on the other hand do not convey egocentric information; seeing a photograph does not tell me anything much about where the object photographed is in relation to me.¹⁴

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.^{15, 16}

The most obvious way of understanding this proposal is as adding a doxastic requirement (a requirement about what the agent believes or knows) to the conditions that an agent must satisfy if she is to count as seeing an object. Understood in this way, the proposal is that seeing requires the formation of certain beliefs or judgments. For example, Currie specifically refers to the “kinds of judgments we make in cases of ordinary seeing. . . which have no counterparts in the case of seeing photographs.”¹⁷ Similarly, Carroll speaks of ordinary seeing as requiring knowledge about spatial relations.¹⁸

Walton has argued that no proposal of this sort can be successful because the requirement it places on seeing is too strong.¹⁹ To make this point, Walton imagines two cases in which a viewer sees a carnation without meeting the doxastic requirement about spatial information set out above.

In the first, a viewer receives visual information about a carnation through a long series of mirrors; the viewer knows neither how many mirrors are involved nor how they are oriented, so he has no idea in what direction the carnation lies from him.²⁰ Walton claims that this viewer will lack information about the location of the carnation in egocentric space; but since all parties to the discussion concede that mirrors are transparent, he thinks, the viewer should count as (prosthetically) seeing the carnation. In the second case, the carnation is indeed right in front of me, but there are many mirrors around, or I suspect that there are. Here, too, Walton claims that I lack the relevant egocentric spatial information about the carnation: “I think I may be seeing the image of a carnation reflected in one or many mirrors. So I have no idea where the carnation is in relation to me.”²¹ Since he thinks that in both cases the viewer sees the

carnation, even though he lacks information about its egocentric location, Walton concludes that possession of that information about the carnation is not necessary for seeing it.²²

While these cases pose serious problems for Currie and Carroll, we do not believe that they settle the issue against the doxastic proposal by themselves. For one thing, although Currie is comfortable denying that seeing takes place in the sequence of mirrors case²³—and would seem forced to take the same position about Walton's second case—an alternative answer would be to weaken the doxastic requirement so as to evade the case. For example, one might hold that seeing requires not (as before) holding a belief about the egocentric location of the object, but merely the belief that the object is in the same general space as oneself. On a weakened doxastic theory of this sort, it is plausible that the agent in both of Walton's cases manages to see, since, plausibly, such very minimal belief is present in these cases. Unfortunately, we anticipate that the debate would become stymied if carried on in this fashion: Walton would respond with further counterexamples to the weakened doxastic requirement, which could then be used to motivate still weaker versions of the doxastic requirement, at which point Walton would concoct yet stranger counterexamples, and so on. We believe that a cycle of counterexamples and responses of this kind is unlikely to convince anyone of anything. However, we propose to sidestep these difficulties: as we shall argue below, there are independent (and, we believe, more compelling) reasons for doubting that any doxastic solution can succeed. It is to these reasons that we now turn.

III. TOWARD A NONDOXASTIC SOLUTION

We are convinced that the contemplated requirement on seeing proposed by Currie and Carroll is too strong. However, we believe that a proper appreciation of the reasons for the failure of this requirement points the way toward a more successful answer to Walton's challenge. Rather than weakening the doxastic requirement, we propose to drop it altogether, while retaining Currie's and Carroll's insights that spatial information is the key to resisting transparency.

The requirement at issue (on the doxastic reading considered so far) concerns what subjects must believe in order to count as seeing an object. Walton's cases are designed to bring out the failure of such a doxastic requirement on object-seeing by pointing out that beliefs can be undermined too easily—namely, beliefs can be undermined in ways that do not undermine seeing. For example, virtually any of my beliefs can be undermined by the onset of a sufficiently far-reaching skepticism. But, while it is plausible that the onset of skeptical doubt might erode a subject's *belief* that she sees a carnation (or her belief that she is within four feet of a carnation, or even her belief that she is somewhere near a carnation), presumably we do not want to say that it would (by itself) prevent her from *seeing* a carnation that is right in front of her face. This is why we are inclined to say of Walton's second case, wherein the subject merely doubts that she lacks egocentric information, that the subject nonetheless sees the carnation. Similarly, the onset of confusion may undermine a subject's *belief* that she sees (or any of her other beliefs, for that matter), but it is implausible that such confusion should (by itself) vitiate her capacity to see. This is why we are inclined to say of Walton's first case, wherein the intervention of a series of mirrors at unknown angles makes the subject confused about the egocentric location of the carnation, that the subject continues to see the carnation.²⁴

These reflections suggest to us that no doxastic condition on object-seeing will suffice to distinguish prosthetic seeing through mirrors from (putative) prosthetic seeing through photographs. That is, object-seeing cannot involve the requirement that the subject believe any particular content, such as content about the egocentric location of particulars. Belief is fragile with respect to perturbations that leave seeing intact, so no doxastic state can be necessary for seeing.²⁵

(A further reason for thinking that a doxastic requirement on object-seeing is too strong involves the possibility of object-seeing by nonhuman animals and neonate human beings. For one thing, while many writers have felt uncomfortable attributing doxastic states to nonhuman animals and human neonates, they have generally been less reluctant to claim that such creatures are incapable of object-seeing; but if object-seeing requires any doxastic state, then the

latter claim follows from the denial of doxastic states in nonhuman animals and human neonates. For another, the question whether all seeing animals are cognizing animals strikes us as broadly empirical; as such, it strikes us as inappropriate as a matter of methodology to allow this question to be settled as a consequence of the requirements on object-seeing imposed from the armchair.)

What then explains the continued appeal of placing a doxastic condition on seeing? Like Fred Dretske, who has also argued for a non-doxastic account of seeing,²⁶ we are inclined to apportion blame to a number of sources. Among these, Dretske points to the “utterance implications” of ordinary statements about seeing. For example, since we do not ordinarily say that we see (or have seen) an object unless we have identified the object that we see, talk about seeing often creates the conversational implicature that we hold identificatory beliefs about the seen object. Dretske also notes that ordinary instances of object-seeing typically lead us to form beliefs about seen objects, and, in particular, we ordinarily form beliefs about their egocentric locations. Finally, we suggest that the appeal of doxastic accounts of object-seeing may be due, in part, to running together the conditions required for seeing, on the one hand, and the conditions that are required for knowing that one sees, on the other. It is plausible that knowing that one sees an object requires that one have beliefs about the egocentric location of the object. But since one can see without knowing that one sees, this observation does not motivate a doxastic condition on mere object-seeing.

IV. EGOCENTRIC SPATIAL INFORMATION

The moral we have drawn so far is that a successful answer to Walton’s challenge cannot involve a doxastic requirement on object-seeing. On the other hand, we do not believe that photographs are transparent, and we are sympathetic to the general idea of exploiting egocentric spatial information to distinguish between genuine and nongenuine cases of prosthetic seeing. Our task, then, will be to find a way of exploiting egocentric spatial information that does not place doxastic requirements on seeing, and thereby avoids the problems that plague the variants examined so far.

The doxastic baggage we are hoping to avoid appeared in other accounts because they understood the notion of (egocentric spatial) information doxastically—namely, in terms of the beliefs or knowledge produced in would-be seers. We propose to do without doxastic elements by construing the notion of (egocentric spatial) information nondoxastically. In particular, we propose to rely here on an understanding of information originally due to Shannon and Weaver, and that was first put to philosophical use (as far as we know) by Dretske in *Knowledge and the Flow of Information*.²⁷

To a first approximation, Dretske understands information-carrying as a kind of (objective) probabilistic, counterfactual-supporting connection between independent variables. Thus, for example, the state of the room’s thermometer carries information about the temperature in the room insofar as there is an objective probabilistic connection between the two: the probability of the room temperature’s being 72°F conditional on the thermometer’s reading 72°F is much higher than the probability of the room temperature’s being 72°F conditional on the thermometer’s not reading 72°F (assuming the thermometer is in good working order, is calibrated, is free from outside influence, and so on). And this probabilistic relationship is made evident by counterfactuals relating the room temperature and the thermometer’s state; namely, if the temperature of the room were different then the thermometer’s reading would be different (again, assuming ideal conditions). Note that there is no such probabilistic link between the state of the thermometer and the temperature of a second room, even if the temperature of this second room is also 72°F; for here the correlation is accidental—not probabilistically linked, and not counterfactual supporting. (Of course, the correlation is not accidental if the temperature in the second room is probabilistically related to the temperature in the first room—say, by thermal connections between the two rooms or something; but in that case there will also be an informational link between its temperature and the state of the thermometer.) Again, evidence for this claim about information relations comes from consideration of counterfactuals; namely, it is not the case that if the temperature in the second room were to change, the thermometer reading would change.²⁸

Significantly, the claim that x carries information about y is a claim about an objective probabilistic link between the two, and as such its truth is independent of anyone's doxastic attitudes about the two.²⁹ In the context of the problems examined in Section III, such a nondoxastic understanding of information—according to which information-carrying is independent of any subject's beliefs about whether it is carried—seems to be just what we need to move forward. We propose to use this nondoxastic account of information to formulate a nondoxastic requirement on object-seeing. We believe that this will enable us to draw the desired distinction between ordinary and prosthetic seeing on the one hand, and photography on the other.

Here, then, is our proposed answer to Walton's challenge. We propose that neither belief nor knowledge about the egocentric spatial location of an object is a necessary condition for seeing it, but instead that what is essential is that the relevant visual experience is produced by a process that carries egocentric spatial information about the object. That is, x sees y through a visual process z only if z carries information about the egocentric location of y with respect to x . According to us, mirrors are transparent in Walton's sense because mirrors carry egocentric spatial information about objects. In contrast, our view secures the desired conclusion that photography is not transparent, insofar as the visual process of looking at photographs fails to carry egocentric spatial information about their depicta. For there is no probabilistic relationship between the photographic image and the egocentric location of the depictum: as I move around the world with the photograph, the egocentric location of the depictum changes, but the photographic image does not.

Some comments are in order.

1. We are not claiming that carrying information about egocentric location of perceived objects is sufficient for seeing. Clearly it is not, and our account reflects this.
2. Information-carrying is, in the first instance, defined for token events: token x carries information about token y just in case there is an appropriate objective probabilistic link between x and y . Our requirement on object perception, in contrast, talks about the information carried by a visual process type—

appropriately, since the challenge we take ourselves to be answering is a challenge to mark distinctions among various process types (not tokens). We intend the appeal to the information carried by process types as a generalization of the more basic informational relation defined for token events. Namely, we can say that a process type carries the information of a certain kind just in case the process's tokens are typically tokens that carry information of that kind. Indeed, we can construe this feature of processes as a disposition: processes that carry information of a certain kind are disposed to have tokens that carry information of that kind.

3. Since, as just noted, the notion of information-carrying for process types is construed dispositionally, we can allow that, even for a process type that carries information about the egocentric location of objects, there may be individual tokens of the type that fail to carry information about the egocentric location of objects represented. That is to say, information-carrying processes can fail to carry information in individual cases without thereby ceasing to be information-carrying processes.
4. Why do we formulate our condition on object-seeing in terms of egocentric location (i.e., location with respect to the viewer) rather than absolute location or allocentric location (i.e., location with respect to a frame of reference independent of the viewer)? Mainly because it enables us to draw the distinctions we need to draw while avoiding a host of thorny issues.

For example, suppose someone wants to individuate photographs by the absolute or allocentric locations of their depicta. Then if the counterfactuals are read *de dicto*, photographs will, trivially, carry information about the allocentric location of their depicta (because the relevant counterfactuals will turn out to be vacuously true). But that would mean that a requirement stated in terms of absolute or allocentric locational information will not distinguish between the visual process of looking at photographs, on the one hand, and uncontroversial cases of prosthetic or nonprosthetic vision on the other. We suppose we could defend an allocentric/absolute formulation of our requirement if we were willing to rule out

the individuating standard at issue or plump for a *de re* reading of the counterfactuals, but we would prefer not to take sides about such tendentious issues if we can avoid it. In contrast (however you read the counterfactuals), individuating photographs by the allocentric location of their depicta does not make it the case that photographs carry information about the egocentric location of their depicta.³⁰ Furthermore, the idea of individuating photographs by the egocentric location of their depicta is implausible on its face—it is hard to accept that walking across the room with the photograph of your grandmother amounts to replacing the old photograph with a new one. Consequently, formulating our requirement in terms of egocentric locational information allows us to distinguish between photographs and visual prosthetics without having to make controversial claims about the individuation of photographs and the interpretation of counterfactuals.

5. Our account does not place a doxastic requirement on seeing. We hold that what is essential to seeing is that the relevant visual experience is produced by a process that carries information about the egocentric spatial information of the perceived object. On our account, knowledge—or even mere belief—about the location of the object is not necessary for seeing. For the activity of information-carrying processes need not result in beliefs (for example, such processes may be at work in the case of the thoroughgoing skeptic, although in her case they would fail to result in beliefs; this is in accord with our insistence in Section III that the onset of skeptical doubt should not erode the capacity for object-seeing). For this reason, our account evades the problem pressed against other answers to Walton in Section III.
6. Walton complains that the doxastic proposal put forth by Currie and Carroll amounts to “ad hoc linguistic legislation,”³¹ and that such disputes over terminology have no philosophical interest. Are we, in putting forward our own anti-transparency proposal, guilty of the same sin?

We hope not. Linguistic usage of ‘sees’ is eclectic, and obviously occurs (appropriately) in cases where object-seeing is not at issue (e.g., when someone says “I see a horse

in the clouds”); cf. Dretske, *Seeing and Knowing*). As such, we agree with Walton that there is no point in arguing over who gets to keep the word. On the other hand, we take it that there is philosophical point to “bring[ing] out the important similarities and differences... especially the kinship which seeing a photograph of something bears to other ways of seeing it, and seeing a painting of it does not.”³² We contend that our own proposal attains this goal more successfully than does the transparency thesis. For while both accounts bring out similarities between object-seeing and the ways in which we perceptually interact with photographs and other depictive representations, our account is more successful than Walton’s in bringing out differences between the cases—differences whose salience and importance is revealed by the extent to which most readers (including Walton himself; see note 11) regard Walton’s proposal as counterintuitive.

To see that our proposal is extensionally correct, let us examine a few cases. We begin by looking at some cases where it is clear that a visual process does underwrite object-seeing.

We take it as fairly obvious that our view allows for ordinary (nonprosthetic) seeing and uncontroversial cases of seeing by visual prostheses. Ordinary seeing carries information about the egocentric location of objects (although it is, of course, subject to failures of information-carrying in individual cases). In addition, our proposal allows for uncontroversial cases of prosthetic vision involving eyeglasses, binoculars, telescopes, and periscopes. All of these prostheses carry information about the egocentric location of objects perceived (although they may fail in certain circumstances). Evidence of this informational link can be found in various counterfactuals that hold true about these processes; in each case, if one were to change the egocentric location of the objects seen, one would be presented with a different visual image.³³ For the same reason, our view allows for seeing through a single mirror. Moreover, the condition we propose creates no problem for saying that I see in the case in which I am surrounded by many mirrors (or merely suspect that I am). Although this situation might undercut my belief that I am seeing, and hence my

ability to know that I see, this cannot undercut the mirror's capacity to carry information about the egocentric location of the carnation. Finally, the same holds true regarding a long sequence of mirrors. In all of these cases, change in an object's egocentric location would bring about change in the (mirror-produced) image.

What about cases in which a visual process does not underwrite object-seeing? As mentioned above, our proposal provides a principled basis for rejecting photographic transparency. That is, it implies that photographs and films do not allow us to see the objects they depict. For, as we have maintained, visual processes involving photographs and film fail to carry egocentric spatial information about their depicta (although they do carry some sorts of information about their depicta): there is no probabilistic relationship between the photographic/film image, on the one hand, and the egocentric location of the depictum, on the other. Evidence of the lack of this probabilistic relation is, as usual, to be found by consideration of the counterfactuals linking the two. It is not the case that if our egocentric location with respect to the objects were to change, the photographic image would change. For one can walk around with a photograph—changing one's location with respect to the depictum—without a concomitant change in the photographic image.

The same seems to be true of broadcast and live-feed video. Broadcast video (whether broadcast live or from a prerecorded source) can be viewed in many different absolute locations; hence the (fixed) depictum can be in any of many different egocentric locations with respect to the viewer without any change in the video image. This shows that broadcast video fails to carry information about the egocentric location of the depictum, hence does not permit prosthetic object-seeing. On the other hand, suppose there is a direct video feed (not a broadcast signal) directly connected to a monitor at only one location. If the monitor should happen to remain in one place, there may very well be a *de facto* correlation between the video image and the egocentric location of the depictum. But this correlation fails to rise to the level of an informational link, as can be seen by consideration of the relevant counterfactual: if, contrary to fact, there were a modification in the egocentric location of the depictum (say, if, contrary to

fact, someone bought a longer video cable and moved the monitor by twenty feet), the video image would remain unchanged. Here, too, then, the process type fails to carry information about the egocentric location of the depictum, so there is no object-seeing.³⁴

What about painting and drawing done in realist (or even photo-realist) style? Consider cases in which a painter strives to depict an actual person or scene accurately. In these cases, counterfactual dependence and the preservation of real similarity relations may be present. In addition, such paintings may carry a great deal of information about their subjects. But visual processes involving these paintings (like photographs) fail to carry egocentric spatial information about the objects they depict. Again, one can move about with the painting—thereby changing one's egocentric location with respect to the objects that it depicts—without changing the image. Hence, seeing the depicted objects is precluded (as desired).

V. THE EPISTEMIC VALUE OF PHOTOGRAPHS

So far we have been reacting against the transparency thesis—we have taken it as our goal to show why photographs are not transparent. However, our account of what separates genuine cases of prosthetic perception from photographs suggests what seems to us a new and better way of addressing one of the original questions to which transparency was proposed as an answer: Why are photographs epistemically special in a way that other sorts of depictive representations are not? (Why, for example, do photographs but not paintings carry evidentiary weight?) This question will be our focus in what follows.

To answer it in a way that connects with what we have said so far about visual process types, we need to extend our information-related terminology once again so that it will apply to depictive representation types. Therefore, we shall say that a type T of depictive representation carries information of a certain kind if and only if the type of visual process involving looking at T-tokens carries information of that kind. Likewise, we can define a notion of information-carrying for token visual representations: a token t carries information of a certain

kind just in case the visual process token involving looking at *t* carries information of that kind. Like the notion of information-carrying defined for visual process types above, the notion of information-carrying for depictive representation types is dispositional; hence, a type *T* can carry information of some kind despite occasional failures to carry information of that kind by isolated *T*-tokens.

With this terminology in hand, we propose that part of what makes photographs epistemically special is this: they are information carriers whose conditions of employment are easier to satisfy than other information carriers. To see what we mean by this, consider these two kinds of information that a representation (or representation-type) can carry about its representational object:

- (i) information about the visually accessible properties of the representational object, and
- (ii) information about the egocentric location of the representational object.³⁵

For many ordinary cases these two sorts of information-carrying capacities coincide: things that carry type (i) information also carry type (ii) information. Even more than this, it seems that things typically have their capacity to carry type (i) information only insofar as they have a concomitant capacity to carry type (ii) information as well. Thus, for example, ordinary, non-prosthetic vision carries both type (i) and type (ii) information, and would not carry type (i) information if it did not also carry type (ii) information. Similarly, a mirror carries type (i) and type (ii) information about the carnation, but it would not carry type (i) information about the carnation if it were not spatially related to the carnation and the viewer in such a way that it also carries type (ii) information.

Ordinarily, then, transmission of type (i) information occurs only when transmission of type (ii) information occurs as well. However, we believe, there are exceptions to this generalization; namely, for reasons given in Section IV, we believe that photographs convey type (i) information without conveying type (ii) information. To coin a term, we shall refer to anything that carries type (i) information without carrying type (ii) information as a *spatially agnostic informant*. We propose to explain the

special epistemic features of photographs (partly) in terms of their being spatially agnostic informants.

Why do we believe that this fact about photographs has interesting epistemic consequences? First, insofar as information about visually accessible properties is epistemically useful, it seems clear that anything that carries information of type (i) has nontrivial epistemic value. Photographs have an epistemic value that paintings and other sorts of depictive representations lack since the former carry type (i) information while the latter do not. (Our claim here is that the criterion in play can distinguish between photographs and paintings *qua* types of depictive representations; we do not mean to deny that there are token paintings that carry information of type (i). More on this point below.)

However, our suggestion is that, insofar as they are spatially agnostic informants, photographs have an epistemic value not possessed by other type (i) information carriers. As noted, most things that carry type (i) information are unavailable except where they also carry type (ii) information. But this is to say that such sources of type (i) information come with strings attached. In particular, they do not serve as sources of type (i) information in those cases where they fail to provide type (ii) information. And, as it happens, sometimes we are not situated in a way that allows for the provision of type (ii) information at all, but where it nonetheless serves our needs to have type (i) information. It is in these cases that photographs hold a special epistemic value. Photographs are epistemically valuable because they constitute a relatively undemanding source of information about the visually accessible properties of objects—one that works even when we lack information about egocentric location.

Insofar as tools with relatively undemanding conditions of employment are valuable, the fact that photographs are spatially agnostic informants explains why they are valuable in ways that even other information carriers are not.³⁶ Of course, this is not to say that photographs are epistemically superior to other sources of type (i) information along every dimension of evaluation; in particular, one dimension along which the former are epistemically inferior to the latter is that, precisely because of their spatial agnosticism, photographs are informationally poorer

than other sources of type (i) information. The special epistemic value of photographs, we claim, is pragmatic: they do their epistemic job of conveying type (i) information in situations where other candidates for the job are unavailable.³⁷

VI. TOKENS, TYPES, AND EVIDENTIARY STATUS

What we have said in Section V is intended to explain why photographs (*qua* type) enjoy an epistemic status that is in some ways superior both to other types of depictive representations (say, paintings) and other types of carriers of type (i) information (say, mirrors). But, of course, some tokens of other representation types—some realistic portrait paintings and some courtroom illustrations, for example—are also spatially agnostic informants;³⁸ thus, if what we have said on behalf of photographs is correct, these nonphotographic tokens should hold the same special epistemic status. On the other hand, this seems counterintuitive, insofar as we do not accord the same epistemic status to realistic portrait paintings that we accord to photographs. What has gone wrong?

The first point to make in connection with this is that, taken as particulars, these tokens pose no threat to the thesis about what makes photographs special: our thesis is intended to mark out what is special about photographs *qua* type, and this is compatible with extending the same status to individual tokens of other types. However, the possibility of gathering these particulars into a type suggests an objection to our view, and a full explanation of the epistemic status of photographs demands an answer to this objection. Consider the type consisting of token depictive representations that are spatially agnostic informants but are not photographs. On our own admission, this type is nonempty. Moreover, there are several interesting and non-empty subtypes of this type, including, for example, the type of veridical portrait paintings. These types are like the type of photographs in that their members are spatially agnostic informants. Yet we take it that there is an important epistemic difference between these types and the type of photographs: it seems that instances of the former types do not command assent in the way that photographs do (for example, they are not accorded the same evi-

dentary weight that is accorded to photographs; but see below for exceptions to this claim). This shows that more remains to be said.

We propose to explain the relevant epistemic difference between photographs, on the one hand, and tokens of (most) other types of spatially agnostic informants, on the other, by holding that the former type is salient for subjects in a sense that these other types are not. By this we mean that subjects visually experiencing a token photograph typically categorize that token as an instance of the type of photographs (typically on the strength of their visual experience of the token and background knowledge about how photographs visually appear, although nothing we shall say hangs on this). In contrast, we claim, subjects visually experiencing a token veridical portrait painting typically do not categorize that token as an instance of the type of veridical portrait paintings; rather, unless they have some special knowledge about the conditions under which the work was made, they are likely to categorize a token of this type as an instance of the type of paintings, or perhaps an instance of the type of portrait paintings.³⁹

Moreover, we contend, subjects (tacitly) hold relevant background beliefs about the types to which they assign these works.⁴⁰ Namely, they believe that the type to which token photographs are assigned (namely, the type of photographs) is a type whose members carry type (i) information. In contrast, they believe that the type to which token veridical portrait paintings are assigned (namely, the type of paintings, or the type of portrait paintings) is a type whose members may fail to carry type (i) information.

These background beliefs, we think, explain why subjects accord a different evidentiary status to token photographs and token portrait paintings. Namely, as a result of the beliefs in question, subjects believe that a currently seen token photograph (like other tokens of the type of photographs) is likely to carry type (i) information. In contrast, subjects will believe that a currently seen token veridical portrait painting (like other tokens of the type of portrait paintings) may fail to carry type (i) information.⁴¹ But, plausibly, a subject will take a depictive representation as evidence for a proposition *p* only if she believes that it (or its type—see point 2 in the itemized list in Section IV) carries the information that *p*. Consequently, subjects

will take photographs as evidence for propositions about the visually accessible properties of their representational objects, but will not extend this evidentiary status to veridical portrait paintings (in the absence of independent knowledge about their veracity).

It is worth noticing, at this point, that what we have said about the type of photographs plausibly extends to some other types of depictive representations as well. For example, it may well be that ornithological illustrations carry information about the visually accessible properties of our avian friends. Moreover, it is plausible that this is a salient type for many typical viewers of such depictions. Their distinctive style, and the contexts in which they are typically presented, make it the case that token ornithological illustrations are typically categorized as tokens of the type of ornithological illustrations rather than simply as tokens of the type of drawings. In addition, it is plausible that subjects typically, (and tacitly) hold the background belief that tokens of this type provide type (i) information. If this is right, then the type of ornithological illustrations (and, potentially, other types as well) would seem to be on all fours with the type of photographs. And this is not implausible; we are prepared to accept that ornithological illustrations carry the sort of evidentiary weight at issue (in the bird-watching and bird-identifying contexts in which the illustrations in question are typically used).^{42, 43}

VII. THE CONTINGENCY OF PHOTOGRAPHIC PECULIARITY

In Sections V through VI we have argued that there are important epistemic differences between the class of photographs and other classes of depictive representations. However, we believe that these differences between photography and other forms of depictive representation are contingent, rooted in the history of the uses of these various media, and not solely in their material natures.

Before we give our reason for believing this claim, we want to consider another reason that might be advanced in favor of the same conclusion, but that we take not to be convincing. The thought here starts from the observation that, on our account, the epistemic peculiarity of photo-

graphs is rooted in their being spatially agnostic informants. But some have alleged that some photograph tokens do, under certain circumstances, provide type (ii) information. For example, Currie writes that “photographs can serve, along with information from other sources, in an inference to egocentric information. If I know where and when the shot was taken, and where I am now (and what the time is now), I may infer that the scene depicted stands in a certain spatiotemporal relation to my current time slice.”⁴⁴ If Currie is right about this, then it follows that spatial agnosticism (and hence whatever epistemic peculiarity of photographs that results from their spatial agnosticism) is a merely contingent feature of photographs. However, this view seems to us to be mistaken. For, at least given our way of thinking about information, photographs do *not* carry type (ii) information in Currie’s cases (despite any inferences that may be drawn from them).⁴⁵ For this reason, we are unpersuaded by the line of thought under consideration.⁴⁶

For all that, we do believe that the epistemic peculiarity of photographs is contingent. To see why, recall that our account of the epistemic differences between photographs and other spatially agnostic depictive representations rests (in part) on claims about the relative salience of various representational types and the standard background beliefs associated with these. We have claimed that the type of photographs is salient in a way that the type of veridical portrait paintings is not (i.e., photograph tokens are typically categorized as photographs, while veridical portrait paintings are categorized merely as paintings or portrait paintings). Moreover, it is plausible that the folk theory of photography assumes that photographs are sources of visual information, while our background beliefs about paintings differ—we do not expect them to be sources of information. But both the saliency ordering among representational types and the generally-held background beliefs about these types are, presumably, contingent; namely, they are contingent on facts about both the history of representational practices and our perceptual/cognitive psychology. That is, if these facts had been different, the evidential status of photographs *vis-à-vis* that of paintings (for example) might have been different as well.⁴⁷

VIII. CONCLUSION

In this paper we have made a number of claims about the epistemic status of photographs. First, we have claimed, *pace* Walton, that photographs are not transparent because, unlike mirrors, telescopes, and the like, they are spatially agnostic informants. However, we suggest, the type of photographs is epistemically superior to the type of drawings in that instances of the former but not the latter type provide type (i) information about their depicta. Moreover, the type of photographs has (and, consequently, tokens of that type have) an evidentiary status superior to the type of portrait paintings (and its instances) and even the type of veridical portrait paintings (and its instances) because of differences in the salience of these types for subjects and their background beliefs about these types. On the other hand, the same sort of factors explain why the special evidentiary status of photographs also extends to some subcategories of veridical nonphotographic depictions in some settings. Finally, we claim, the epistemic differences between photography and other depictive representations are based on contingent rather than necessary features of these representational types.⁴⁸

JONATHAN COHEN

Department of Philosophy
University of California, San Diego
La Jolla, California 92093

INTERNET: joncohen@aardvark.ucsd.edu

AARON MESKIN

Department of Philosophy
Texas Tech University
Lubbock, Texas 79409

INTERNET: aaron.meskin@ttu.edu

1. Gerhard Richter: *The Daily Practice of Painting: Writings and Interviews 1962–1993*, ed. H. U. Obrist (MIT Press, 1995).

2. Some of the material in this paper appeared (in an earlier version) in a shorter paper titled “Photographs Are Not Transparent” that we presented at the 2003 Pacific Division meeting of the American Society for Aesthetics. This work is fully collaborative; the authors are listed alphabetically.

3. A. Bazin, *What Is Cinema?* trans. Hugh Gray (University of California Press, 1967), p. 14.

4. Kendall Walton, “Transparent Pictures: On the Nature of Photographs,” *Critical Inquiry* 11 (1984): 246–276.

5. To be fair, Walton uses the transparency thesis to explain more than just the epistemic value of photographs. That said, this is clearly one of the important explanatory targets that he uses to motivate the proposal.

6. Cinematographic and video depictions also count as transparent on Walton’s account; the intended contrast is with painting and drawing, which he takes to be nontransparent. Note that Walton does not claim that the depictum is the only thing we see when we look at a photograph; in particular, he does not deny that we see the photograph in addition to its depictum. Indeed, he insists that it is in virtue of seeing the photograph that we see its depictum. Hence, on this view, transparency does not entail invisibility.

7. Compare David Lewis, “Veridical Hallucination and Prosthetic Vision,” *Australasian Journal of Philosophy* 58 (1980): 273–286.

8. Note that photographs are not unique among these visual prostheses in allowing for a specifically temporal separation between viewer and the object seen: we speak unhesitatingly of seeing a stellar explosion through a telescope, even if the explosion transpired millions of years before the viewer existed.

9. Walton, “Transparent Pictures: On the Nature of Photographic Realism,” pp. 251–252; emphasis in original.

10. Gregory Currie, *Image and Mind: Film, Philosophy, and Cognitive Science* (Cambridge University Press, 1995), p. 55.

11. We suspect even Walton would concede this much—this would explain why he felt the need to warn against taking the thesis nonliterally. Of course, these sorts of intuitions are not infallible, so it is open to one to respond by rejecting them. On the other hand, general canons of rational conservatism counsel against rejecting such intuitions (especially deeply held ones) when less revisionary alternatives are available. Therefore, we propose to take the intuitions seriously and attempt to explain them as reasonably as we can. This policy applies to both the intuition that photographs provide a special epistemic connection to the world (one that Walton accepts) and the anti-transparency intuition (one that Walton rejects even while appearing to recognize its force), *inter alia*.

12. Noël Carroll, “Towards an Ontology of the Moving Image,” in *Philosophy and Film*, ed. C. A. Freeland and T. E. Wartenberg (New York: Routledge, 1995), pp. 68–85 and *Theorizing the Moving Image* (New York: Cambridge University Press, 1996), pp. 62–63; Gregory Currie, “Photography, Painting, and Perception,” *The Journal of Aesthetics and Art Criticism* 49 (1991): 23–29, and *Image and Mind: Film, Philosophy, and Cognitive Science*; and Nigel Warburton, “Seeing through ‘Seeing Through Photographs,’” *Ratio* 1 (1988): 64–74.

13. Some articulations of this point (e.g., that in Currie, *Image and Mind: Film, Philosophy, and Cognitive Science*, p. 66) put the point in terms of spatiotemporal relations. We prefer to express the point in terms of spatial relations in the context of an attempt to exclude photographs because, arguably, when *x* looks at time *t* at a photograph of *y*, *x* (or *x*’s visual system) represents the information that *y* existed before time *t*. We do not see any nonstipulative reason for refusing to count this as information about the spatiotemporal relation between the viewer and the depictum, but it does

not seem reasonable to count it as information about the spatial relation between the viewer and the depictum.

14. Currie, *Image and Mind: Film, Philosophy, and Cognitive Science*, p. 66.

15. Carroll, *Theorizing the Moving Image*, p. 62.

16. Currie and Carroll claim that the spatial requirement in question is a necessary condition for prosthetic seeing, not that it is a sufficient condition. This is all to the good, since it is not a sufficient condition: if I am looking straight down at my desk, wearing blinders, and you hand me written descriptions of the spatial relations that obtain between me and objects in my vicinity, then I may know where these objects are in relation to me, but presumably I am not (or not literally) seeing these objects prosthetically (using you as my prosthetic).

17. Currie, *Image and Mind: Film, Philosophy, and Cognitive Science*, p. 66.

18. We find the doxastic construal the most straightforward reading of Currie and Carroll. We shall be arguing below that no such doxastic proposal can succeed as an answer to Walton, and offering a nondoxastic proposal in its stead. However, if Currie and Carroll want to insist that they originally intended a nondoxastic view, and that our proposal is a mere extension of what they had in mind all along, that's all right with us too.

19. Kendall Walton, "On Pictures and Photographs: Objections Answered," in *Film Theory and Philosophy*, ed. R. Allen and M. Smith (Oxford University Press, 1997), pp. 60–75.

20. Walton, "On Pictures and Photographs: Objections Answered," p. 70.

21. *Ibid.*

22. We can imagine a defender of the egocentric information requirement who would allow that, after the number of intervening mirrors between the subject and the carnation gets sufficiently large—say, greater than n , the subject ceases to see the carnation. Therefore, she might suggest, the case involving $n+1$ mirrors is not a case where the subject sees without egocentric spatial information, hence not a counterexample to the requirement she is defending. But we find this response unconvincing. For as Walton's second case shows, the point does not turn on assuming large numbers of mirrors are involved. Therefore, the point goes through even if we concede the objection.

23. Currie, *Image and Mind: Film, Philosophy, and Cognitive Science*, p. 70.

24. The intuitions about Walton's specific cases adduced here are certainly not beyond dispute, especially if the doxastic requirement under discussion is weakened in the way imagined at the end of Section II. But the general moral we are drawing stands, independently of verdicts about these specific cases: mere confusion can undermine belief but cannot undermine seeing.

25. Arguably there is a nondoxastic reading of at least Carroll's version of the spatiotemporal information proposal. For, at times, Carroll seems to be suggesting that the relevant difference between ordinary seeing and photographic looking has to do with their relation to our physical abilities:

I can "orient my body" spatially to what I see, either with the naked eye or through a telescope or microscope. But when I see a photograph I cannot orient my body to the photographed objects. The space of the objects is "disconnected phenomenologically from the space I live

in." (Carroll, "Towards an Ontology of the Moving Image," p. 71)

If the "orientability requirement" Carroll suggests here is understood as not placing doxastic requirements on would-be seers, then it would evade the problem we have been discussing. However, this requirement, too, seems too strong, since it would inappropriately follow from the requirement (on its most straightforward reading) that organisms incapable of moving their bodies (e.g., normal human victims of paralysis) cannot see any objects.

26. Fred I. Dretske, *Seeing and Knowing* (Chicago: University of Chicago Press, 1969). See also Fred I. Dretske, "Simple Seeing," in *Body, Mind, and Method: Essays in Honor of Virgil C. Aldrich*, ed. D. F. Gustafson and B. L. Tapscott (Boston: Reidel, 1979), pp. 1–15. Indeed, in these works Dretske makes points similar to those we are urging against doxastic accounts. However, the view of nonepistemic seeing defended in *Seeing and Knowing* (according to which "S sees_n D=D is visually differentiated from its immediate environment by S" [p. 20]) plausibly entails that we see objects by means of photographs of them. If so, then Dretske's view also underwrites the transparency thesis. Needless to say, we do not accept this (nondoxastic) account of object-seeing either.

27. Fred I. Dretske, *Knowledge and the Flow of Information* (MIT Press, 1981).

28. It is worth mentioning, only so that we can set aside, two concerns about the metaphysics of information. First, as Barry Loewer (Loewer, "From Information to Intentionality," *Synthese* 70 [1987]: 287–317) points out, it is not obvious how to understand the objective conditional probabilities (including objective conditional probabilities for unrepeatably token events) undergirding Dretske's notion of information: standard proposals are inapplicable, and Dretske does not offer anything in their place. We have nothing to say about this problem. Second, in appealing to certain counterfactuals as evidence for informational relations, we are not committed to the view that informational relations are constituted by the holding of such counterfactuals *per se*. We are only committed to the weaker claim that the counterfactuals provide evidence for the informational relations.

29. Famously, Dretske appeals to this understanding of information, in part, as part of an account of knowledge and justification. Since information-carrying is nondoxastic in the sense explained, it turns out that, on Dretske's account, a subject can know that p (be justified in believing that p) so long as she satisfies the conditions (including informational conditions) *even if she fails to believe that she satisfies them*. In the standard jargon, this means that Dretske's account of knowledge and justification is externalist rather than internalist (cf. L. Bonjour, "Externalism/Internalism," in *A Companion to Epistemology*, ed. J. Dancy and E. Sosa [Blackwell, 1992], p. 132). In helping ourselves to the Dretskean understanding of information in the service of a proposal about object-seeing, we are not thereby signing on to Dretske's epistemology, and therefore do not take on the burdens of defending Dretske's conception of knowledge or justification.

30. Reading the counterfactual *de dicto*, it is not the case that if the depictum were to change its egocentric location then the image would change, because such a change is only possible if the viewer moves. On a *de re* reading, the

counterfactual fails as well: it is not the case that if Grandma were to change her egocentric location then the image would change.

31. Walton, "On Pictures and Photographs: Objections Answered," p. 71.

32. *Ibid.*

33. Indeed, it is critical to the standard use of these tools that they carry information about egocentric location. Consider, for example, binoculars and periscopes: we do not simply want to find out what the enemy soldier and enemy battleship look like, we also want to know where they are in relation to us.

34. Objection: What was just said about live-feed video plausibly extends to telescopes (etc.), contrary to our claim that these are visual prostheses. For while there is a *de facto* correlation between the telescopic image and the egocentric location of the depictum, one could modify the telescope—by inserting a tube, adding mirrors so the eyepiece is three feet to the left, or whatever—without thereby changing the telescopic image.

Response: First, many changes in telescopic mirrors, tubes, and the like do change the telescopic image. Adding tubes, for example, changes the magnification, while adding a mirror inverts the image. Of course, such changes can be offset by various means (adding mirrors only in pairs so that inversions are reinverted, changing the lenses to offset magnification differences). If you make such extensive modifications to your telescope, it will be true that the egocentric location of the depictum with respect to the viewer before the modifications is distinct from the egocentric location of the depictum with respect to the viewer after the modifications, although the telescopic image will be unchanged. However, we are inclined to say that the telescope before the changes is distinct from the telescope after the changes (and that the type of visual processes involving the first is distinct from the type of visual processes involving the second). Moreover, it seems that there is a probabilistic link between the telescopic image produced by each telescope and the egocentric location of the depictum with respect to it. We claim, then, there is no single telescope such that visual processes involving it fail to carry egocentric spatial information; instead, there are two telescopes such that visual processes involving each of them do carry egocentric spatial information. That is just what one would expect if, as we claim, the type of visual processes involving telescopes carry egocentric spatial information.

35. Objection: If egocentric location is a visually accessible property, then type (ii) information is a species of type (i) information.

Response: In this case, restrict type (i) information so that it includes only information about visually accessible properties of the representational object other than its egocentric location.

In what follows, we shall ignore this point for ease of expression.

36. The importance that photographs have depends partly on the (presumably contingent) fact that we sometimes want type (i) information in situations where type (ii) information is unavailable, given the *de facto* limitations on our perceptual capacities. By way of contrast, notice that black and white photographs are less demanding sources of visual information than color photographs in that the former are chromatically agnostic informants. But since we are rarely confronted with situations in which seeing is possible

but seeing in color is not, we tend not to accord to black and white photographs an epistemic value over and above that of color photographs.

37. A number of authors have held that the special epistemic status of photographs is to be explained (at least partly) in terms of the automaticity, or lack of intentional mediation in the production of photographs. Have we erred in ignoring this factor in our account? We think not; photographic production is indeed mediated by intentions (for example, intentions at work in the selection of lenses, and in the development process; cf. Joel Snyder and Neil Walsh Allen, "Photography, Vision, and Representation," *Critical Inquiry* 2 [1975]: 143–169). On the other hand, it may be that widespread belief in the lack of intentional mediation in photography partially explains the fact that subjects tend to believe that photographs carry type (i) information. If so, this would be relevant to the assignment of high evidential status to the type; we consider this topic in Section VI.

38. It follows from this, of course, that Richter is wrong in saying (see the epigraph at the beginning of this essay) that the photograph is the only picture that can convey information.

39. Why does a subject categorize a given token as an instance of type T rather than type T' (assuming the token is in fact an instance of both types)? We suppose this has to do, in large part, with the degree to which the alternative types are culturally entrenched, the subject's personal history of exposure to the alternative types, and so on. Compare Walton, "Categories of Art," *Philosophical Review* 79 (1970): 357ff. More on these themes in Section VII.

40. Note that, by appealing to belief here we are not giving up our claim that seeing is independent of belief (see Section III); instead, we are appealing to belief only in order to mark out different ways in which subjects tend to treat different types of depictive representations.

41. Significantly, these differing assessments of the information-carrying credentials of photographs and portrait paintings may be erroneous in individual cases; for example, as noted, there are token portrait paintings that carry type (i) information perfectly well. That is to say, a token may in fact be a spatially agnostic informant; but if this fact is unknown to its viewers, it will not be accorded evidentiary weight. Likewise, assuming that altered photographs count as photographs (something that is not obvious, given what we have said in this paper), a token photograph may fail to carry type (i) information; if so, then our account predicts that it will be accorded an evidentiary status that it does not deserve. Evidence concerning another sort of error in judgment provides further support for our view. Namely, it seems that what drives the evidentiary status I accord to a work is not what type of work it is, but rather what type of work I take it to be: if I mistake a photograph for a painting, I will give it less evidentiary weight than other representations I take to be photographs (even if, in fact, the latter are paintings I have mistaken for photographs). We take this as a confirmation of our strategy of explaining the differing evidentiary status of photographs and paintings in terms of subjects' category judgments about works.

42. In other contexts, namely those in which the relevant background beliefs are not in place, our account predicts that the very same tokens will lack the relevant sort of evidentiary weight. Indeed, we can imagine certain cognitive

purposes for which ornithological illustrations (for example) may be epistemically preferable to photographs. For example, suppose that one is interested primarily in making fine distinctions among visually similar birds (as opposed to getting accurate information about the visually accessible features of birds). In this case, it may be preferable to consult representations that highlight and exaggerate dissimilarities between species of birds, even if this highlighting and exaggerating has the result that such representations carry less type (i) information about their depicta than photographs would.

43. What we have said here leaves room for further significant differences between photographs and veridical drawings. Among these, we note that while photographs and veridical drawings are both epistemic tools with relatively undemanding conditions of use, photographs are in one way even less demanding than the latter. This is because extant veridical drawings—although they are spatially agnostic themselves—are necessarily causally dependent on spatially committed informants (viz., the object-seeing of the depictum by the agent who produced the drawing). In contrast, not all photographs depend on such spatially committed informants (despite being intentionally mediated—see note 37). Although many photographs are produced by agents who see the depictum, cameras may also be positioned to photograph automatically in circumstances where no agent sees the depictum (and, hence, where there are no spatially committed informants). For this reason, photographs can convey visual information about places too distant or inaccessible for the use of nonprosthetic or prosthetic vision. This confers upon photography another pragmatic advantage (analogous to the advantage it has over mirrors by being spatially agnostic): photographs can carry visual information in situations where we want that information, and where other candidates are unavailable for the job.

We also hold that, if the category of veridical machine-made drawings were to become salient, their epistemic value would be equivalent to that of photographs, and that they would share the pragmatic advantage just ascribed to photographs.

44. Currie, *Image and Mind: Film, Philosophy, and Cognitive Science*, p. 66.

45. To the extent that Currie would deny this contention, this gives further support to our interpretation of him as holding a doxastic understanding of spatial information (see note 18).

46. In fact, the cases we can imagine in which photographs carry type (ii) information (in our sense) are pretty far-fetched—maybe even (nominally or metaphysically) impossible, for all we know. Still, we have no proof of the impossibility of photographs providing type (ii) information; and if there could be such a case, this would show that our account of the epistemic status of photographs and the basis of our denial of the transparency thesis rely on merely contingent features of photographs.

In order to avoid the aforementioned conclusion about our argument against transparency (which might suggest the possibility of transparent photographs in distant possible worlds), one could respond by amending the necessary condition on object-seeing to rule out processes that are *possibly* spatially agnostic. On such a revised view, the proposed necessary condition on object-seeing—one satisfied by visual prostheses but not by photographs—would be this:

x sees y through a visual process z only if z necessarily carries information about the egocentric location of y with respect to x.

On this account, even if there are possible worlds in which photographs provide information of type (ii), they are still not worlds in which photographs are transparent. Nonetheless, since photographs would not be spatially agnostic informants in such worlds, our view predicts that the epistemic status of photographs in these worlds would be different than that of photographs in the actual world.

47. In this connection, it may be pertinent to consider actual developments in the practice and use of photography rather than (or in addition to) mere possibilities. While it has long been possible physically to manipulate photographs (e.g., through manual or chemical means), thereby degrading the information that they provide, standard photographic practice has eschewed such manipulations and hence photographs have remained a source of information about the visually accessible properties of their depicta. This (contingent) norm, we claim, explains the persistence of our background beliefs about (hence the persistence of the evidentiary status of) photographs. However, the development of digital photography, and, in particular, the development of cheap and easy digital means of manipulating photography, may force viewers to confront seriously the possibility of unreliable, noninformation-carrying photography; if so, this may change epistemological attitudes toward photographs.

As Barbara Savedoff puts the point, “To the extent that we can see photographs as potentially indistinguishable from their digitally altered counterparts, they become suspect as carriers of even the most basic information, suspect as bearers of any evidence” (*Transforming Images* [Cornell University Press, 2000], p. 201). In fact, Savedoff suggests that this shift in status may not even await the development of digital photographs that are perceptually indistinguishable from traditional photography: “If we reach the point where photographs are as commonly digitized and altered as not, our faith in the credibility of photography will inevitably, if slowly and painfully weaken, and one of the major differences in our conceptions of paintings and photographs could all but disappear” (p. 202).

48. We are grateful to Colin Allen, Craig Callender, Dom Lopes, Ram Neta, Jenefer Robinson, Rob Rupert, Scott Walden, and Jonathan Weinberg for helpful comments on earlier drafts.