Photographs as Evidence*

Aaron Meskin† and Jonathan Cohen‡

We cannot conceive of a more impartial and truthful witness than the sun, as its light stamps and seals the similitude of the wound on the photograph put before the jury; it would be more accurate than the memory of witnesses, and as the object of all evidence is to show truth, why should not this dumb witness show it?

— Franklin v. State of Georgia, 69 Ga. 36; 1882 Ga.

1 Introduction

Photographs furnish evidence. This is true in both formal and informal contexts. The use of photographs as legal evidence goes back to the very earliest days of photography, and they have been used in American trials since around the time of the Civil War. Photographs may also serve as historical evidence (for example, about the Civil War). And they serve in informal contexts as evidence about all sorts of things, such as what we and our loved ones looked like in the past.

Photographs are not, of course, the only sorts of pictures that can furnish evidence. It is not hard to think of cases in which a non-photographic picture might serve as legal evidence (perhaps of art theft?). Cave paintings serve as anthropological evidence of human habitation. And crude pictures chalked on the side of a building may serve informally as evidence of adolescent mischief. Nonetheless, photographs seem to have a distinctive epistemic status as compared with other sorts of pictures. Unlike the aforementioned examples of non-photographic depictions, photographs typically provide evidence specially about what they depict. Most significantly, the epistemically special character of photographs is revealed by this fact: we are inclined to trust them in a way that we are not inclined to trust even the most accurate of drawings and paintings.

The traditional explanation for the special epistemic status of photography is that photography is an inherently realistic medium. This certainly sounds

*This work is fully collaborative; the authors are listed in reverse alphabetical order.

†Department of Philosophy, Texas Tech University, Box 43092, Lubbock, TX 79409, aaron.meskin@ttu.edu

‡Department of Philosophy, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0119, joncohen@aardvark.ucsd.edu

right, but it is hard to evaluate the claim without offering a theory of what photographic realism consists in. Unfortunately, a number of the traditional accounts of photographic realism that purport to explain photography's special epistemic status rest on rather unpalatable commitments. For example, Andre Bazin's re-presentationalist theory [Bazin, 1967] makes the implausible claim that "the photographic image is the object itself." And Kendall Walton's claim that photographs are "transparent" [Walton, 1984] — in that they literally allow us to see the objects that they depict — is both counterintuitive and (as we have argued in a previous paper) simply mistaken [Cohen and Meskin, 2004]. Of course, there are also a number of less metaphysically questionable theories of photographic realism. For example, Gregory Currie offers a sophisticated similarity account of depictive and cinematic realism in his [Currie, 1995]. But this account offers little hope of explaining the epistemic status of photography, in part because it makes no distinction between depictive realism in general and photographic realism in particular.[1] The epistemic status of photography, then, has gone either unexplained or misexplained; either way, the application of the notion of realism serves merely to label some of the distinctive features of photographic representation that need explaining. (We think much the same is true of claims about the objectivity of photographic representation.) In the rest of this paper, we offer an explanation of the epistemic status of photography — an explanation of the sort that has either gone missing or come with implausible metaphysical costs in alternative accounts.

2 The Epistemic Status of Photography

What, then, does the distinctive epistemic status of photography consist in? On our view, there are two features of photography that underwrite its distinctive epistemic value: (1) token photographs are spatially agnostic sources of information, and (2) viewers hold background beliefs about the category of photographs that influence their attitudes towards the epistemic status of viewed token photographs. We'll unpack these claims in turn.

What do we mean by saying that photography is a spatially agnostic source of information? In order to explain this, we will first have to say a bit about what we mean by 'information'.

Our view of information is inspired by the account initially presented by Fred Dretske in his *Knowledge and the Flow of Information* [Dretske, 1981]. For Dretske, information is carried when there is an objective, probabilistic, counterfactual-supporting link between two independent events.[2] So, for example, the state of a thermometer carries information about body temperature just in case there is an objective probabilistic connection between the two: the prob-

[1] Nelson Goodman's conventionalist theory [Goodman, 1976] is another view of pictorial realism that offers little hope of explaining the distinctive epistemic status of photography.

[2] We do not want to commit ourselves to all the details of Dretske's development of the notion of information. In particular, we do not commit ourselves to his requirement that the probabilities in question turn out to be unity.

ability of the temperature being 100° conditional on the thermometer reading 100° is much higher than the probability of the temperature being 100° conditional on the thermometer not reading 100° (subject to some provisos of course). Furthermore, this probabilistic relationship must be counterfactual-supporting. In order for there to be an information-carrying relation between the thermometer's state and a person's temperature it must be the case that (ceteris paribus) if the temperature were different, then the thermometer would read differently. In fact, we believe that the presence of this counterfactual relationship is generally excellent evidence for its existence.[3]

It is our contention that photographs are, in this sense, a significant source of information. In particular, photographs typically provide information about many of the visually detectable properties of the objects they depict. Consider shape and size properties. While photographs often distort and deceive with respect to this class of features, it is our contention that they nonetheless typically carry information about them. For example, if the size or shape of a depicted flower had differed then the photographic image of it would have been different. (A bit more precisely, if the size or shape of depicted features of the flower had differed then the photographic image would have been different. Photographs do not carry information about all the visually detectable properties of the objects they depict. For example, they do not typically carry information about the hidden sides of depicted objects.) Another example of this information-carrying capacity can be seen in color photographs. These photographs typically carry information about the color of the objects they depict — if the colors of the objects had been different then the photographic image would have been different. Photographs also often carry information about many of the visual aesthetic properties of the objects they depict (hence their utility in picking travel destinations and prospective dates).[4]

It is worth emphasizing that, on our construal (viz., Dretske's construal) informational links are constituted independently of any subject's beliefs or other mental states. The state of the thermometer carries the information about the temperature of the room whether or not I (or anyone else) infer from the former to the latter, whether or not I am (or anyone else is) aware of the former, and so on. Likewise, questions about the information carried and not carried by photographs (etc.) are to be answered by considering the objective probabilistic relations they bear to various events. In particular, such questions are *not* to be answered by asking what any subjects are justified in believing on the basis of photographs, what they learn about the world from photographs, what they

[3]But not always — in particular, there are cases where the counterfactuals in question will come out vacuously true, but where we want to say that information is not carried. We'll consider some of these cases below.

[4]More precisely, photographs carry information about the visually detectable properties of their depicta *at a time* — viz., at the time the photograph was taken. For suppose on Tuesday I paint (i.e., apply paint to) the flower I photographed on Monday; I have thereby changed the color of the flower without producing a corresponding change in the photograph. That is, the photograph on Tuesday fails to carry information about the color of the flower on Tuesday. On the other hand, the photograph on Tuesday continues to carry information about the color of the flower on Monday. We'll ignore this complication in what follows.

are prepared to infer from photographs, and the like.

Furthermore, this account of the information carrying capacity of photographs makes no reference to the realism or objectivity of photographs, nor to their accuracy, nor (as should be clear from the prior paragraph) does it imply anything about whether or not we ordinarily make correct judgments on the basis of photographs.[5] It is also worth noting that this account allows that two photographic images that carry the same information might look quite different (perhaps because of differences in photographic processing). As Goodman notes in a similar context, systematically replacing the colors of a picture with their complements would not thereby change the informational content of that picture [Goodman, 1976].

But while photographs typically carry information about many of the visual properties of the objects they depict (call this 'v-information'), there is another category of information that photographs fail to provide: information about the egocentric spatial location of the objects that they depict (call this 'e-information'). That is, photographs do not typically provide information about the location — with respect to viewers of that photograph — of the objects they depict. The lack of this information is made evident by the falsity of the relevant counterfactuals. Consider again that aforementioned photograph of the flower. It is not the case that if the spatial relationship between the photograph and the flower were to change that the image of the flower would change. Hence, a viewer of the photograph who carries it around with her could change her spatial relationship to the flower without any change in the photographic image. And so, the photograph (and the visual process involving looking at the photograph) fail to carry e-information about the objects it depicts.

Up to this point, we have primarily focused on the information carrying capacity of token photographs. For ease of exposition, it may be useful to talk of the information carrying capacities of various depictive types, including the category of photographs. We will say that a depictive type D carries information of kind K just in case tokens of type D typically carry K-information. Alternatively, we may characterize this feature of types dispositionally: D carries K-information if and only if it is disposed to have tokens that carry K-information. It will also be useful to apply the notion of information carrying to various visual process tokens and types.[6] Thus, let us say that a visual process token carries information of a certain kind about an object just in case there is an objective probabilistic relation between the process token and the relevant features of the object. Then the visual process token v of looking at a depictive token d carries K-information about an object o depicted in d if and only if there is an objective probabilistic relation between v and the K-features of o. Visual process types carry K-information if and only if their tokens are typically such that they carry

[5] Of course, if someone should choose to use terms such as 'realism' or 'objectivity' to characterize the information carrying capacity of photography, we would have no quarrel with this choice of labels.

[6] The formulations here depart slightly from those in [Cohen and Meskin, 2004]; we've made these modifications in order to evade some (rather technical) worries applicable to the earlier formulations.

K-information. Accordingly, the visual process type of looking at a depictive type D carries K-information about the objects depicted in tokens of D just in case tokens of that process type typically carry K-information about objects depicted by tokens of D.

If this analysis is right, then photographs are what we have called in an earlier paper 'spatially agnostic informants': they carry information of one sort (v-information) while failing to carry information about the egocentric locations of the objects they depict (i.e., e-information). And this begins to explain the distinctive epistemic value of photography. In the first place, photographs are significant source of information. In this way, the type of photographs differ from many other depictive representation types that do not carry information about the visually accessible properties of the objects they depict. In the second place, the spatial agnosticism of the type of photographs, and therefore the process of looking at photographs, differentiates photographs from some other sources of visual information. For example, both ordinary vision and visual processes that use visual prosthetics such as telescopes and binoculars (hereafter 'visual prosthetic processes') carry both v-information and e-information. This might make it seem that the process of looking at photographs is informationally impoverished in comparison with these other sources of visual information. This is correct. Ordinary vision does typically provide more information than does the process of looking at photographs. And visual prosthetic processes are also informationally richer in this way than the process of looking at photographs (i.e., the former provide both v-information and e-information about the objects they allow us to see, while the latter only provides v-information about depicted objects). Nonetheless, this feature also contributes to the special epistemic status of photography. For it is not simply that ordinary vision provides both v-information and e-information; it would not provide v-information unless it also provided e-information. The same is true of visual prosthetic processes. But photographs and the process of looking at photographs have the capacity to provide v-information without providing e-information. This makes photography a particularly valuable technology.[7]

Many things that carry v-information also carry e-information. These sources of v-information come with strings attached. They do not provide v-information in contexts in which e-information is unavailable. Yet we often find ourselves in contexts in which we want v-information while the preconditions for the acquisition of e-information cannot be satisfied. In such cases photographs display their distinctive epistemic value. They are a relatively undemanding source of information about the visual properties of objects — a source that is available in contexts in which e-information is not.

While we believe that photography's status as a spatially agnostic source of information is a large part of the story, this cannot be the full story about the epistemic value of photography. For there are certainly other token depictions that share this feature with photographs. It is plausible that many land-

[7]Of course, the process of looking at photographs does provide e-information about those very photographs. But it is e-information about photographically depicted objects that is at issue. Photographs don't ordinarily depict themselves.

scape paintings and still lifes carry information about the visually detectable properties of the objects they depict. Moreover, these pictures fail, just like photographs, to provide e-information. So these pictures are also spatially agnostic sources of information. Furthermore, it is possible to group some of these tokens together to make non-empty categories of spatially agnostic informants (e.g., consider the category of veridical still lifes). Yet we typically treat these depictions differently than we treat photographs. We do not typically accord the same epistemic status to still lifes and landscape paintings that we do to photographs.

To explain this difference, we propose to appeal to a further component of our view, which involves psychological facts about viewers. First, we claim that the type of photographs is salient for subjects in a way that the type of veridical landscape paintings are not. Subjects who come into visual contact with a photograph (under ordinary viewing conditions) typically categorize that object as a photograph; in contrast, subjects who come into visual contact with veridical landscape paintings do not typically categorize them as veridical landscape paintings, but rather as paintings or landscape paintings (or, perhaps, Dutch landscape paintings). Second, it is plausible that viewers typically hold background beliefs about the depictive categories to which they assign these pictures. By and large, viewers believe that the type of photographs is one whose members carry v-information. And by and large, viewers believe that the categories to which they assign veridical paintings are ones whose members may fail to carry v-information. It is these background beliefs that explain the epistemic distinction we make between photographs and landscape paintings.[8]

Our proposal explains the epistemic status of photographs in this sense: it lays out the facts about photographs and the background beliefs that people hold about the photographs they see, and shows how these facts leads subjects to treat the photographs they see as carrying evidential weight. In contrast, our account does not (and is not intended to) explain the epistemic status of photographs in the sense that it *justifies* or *provides warrant for* the evidential weight they are accorded in all cases. On the contrary, we think that some of the relevant background beliefs are false in some cases, hence that the token photographs involved in these cases don't deserve the epistemic weight they are given. (That said, we think the relevant background beliefs are true in many cases, hence that the token photographs involved in these cases do deserve the epistemic weight they are given.) Still, we think we make the extensionally correct predictions: subjects do indeed accord this (sometimes justified/warranted, sometimes unjustified/unwarranted) evidential weight to the photographs they see. Our account of the epistemic status of photographs, then, is an account of that in virtue of which photographs have the epistemic status they in fact have; it is not an attempt to say that this epistemic status is (always) deserved or

[8]The explanation we are offering cites two contingent features of human psychology: (i) the saliency ordering among spatially agnostic representational types, and (ii) the background beliefs cited in such explanations. For this reason, we are committed to saying that the epistemic distinction between photographs and landscape paintings is contingent, and also a result of relational (mind-involving) facts about those types.

justified.

3 Clarifications, Objections, and Replies

While, we believe, our account resolves many outstanding puzzles about the epistemic status of photographs, it also raises a number of interesting questions of its own — questions about the details of how the view is to be understood, and questions about its applications to some difficult cases. We'll attempt to answer such questions in this section.

3.1 Object Seeing and E-Information

In our view, spatial agnosticism is important not only for the purpose of understanding the epistemic status of photographs, but also because it can be used to draw a line between sources of v-information whose employment constitutes (object) seeing, on the one hand, and sources of v-information whose employment does not amount to (object) seeing. In particular, we contend that it is a necessary condition on a process that counts as object seeing that it carry e-information. This classification groups together processes of ordinary seeing (by means of a normal visual system), prosthetic seeing through eyeglasses, prosthetic seeing through binoculars, etc., as instances of object seeing. In contrast, we contend, looking at a photograph may be a good way of visually acquiring true beliefs about your grandmother, but it is not a way of seeing your grandmother (but see §3.2 for discussion of some reasons why photographs may fail to provide us with true beliefs). However, some commentators have objected that our condition on object seeing is either too strong, or too weak (or both): too strong, because it seems wrongly to preclude seeing by Balint's syndrome patients (for example), and too weak because it inappropriately classifies video with binoculars and other visual prosthetics rather than with photographs. We think both of these objections are incorrect. Let us say why.

Consider the video cases first. The thought here is that our account inappropriately treats video, unlike photographs, as a visual prosthetic; but, insofar as video is something like a temporally extended sequence of photographs, an acceptable account should classify video and photographs together. Our response to this objection will be to say that, while there are different sorts of video that ought to be considered separately, our view delivers the desired verdict that all these sorts of video are (like photographs) spatially agnostic informants, hence that they do not facilitate object seeing.

First consider video that is either broadcast live or from a pre-recorded source. If a video signal is broadcast then it can be viewed by a (suitably equipped) perceiver in many different allocentric locations; hence the (fixed) depictum of the video image can be in any of many different egocentric locations with respect to the viewer without any change in the video image. This shows that (live or pre-recorded) broadcast video fails to carry information about the egocentric location of the depictum, and so counts as another spatially agnostic

informant. Cases of non-broadcast video are slightly more complicated. Some security systems, for example, involve a direct video feed from a single stationary camera to a single stationary monitor. In such cases, it might happen that the camera, the monitor, and the viewer all remain fixed in allocentric space. If so, then there is a *de facto* correlation between the video image and the egocentric location of the depictum. But this *de facto* correlation doesn't make for an informational (i.e., counterfactual-supporting) link: 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; hence, non-broadcast video is also a spatially agnostic informant.

Let us turn now to Balint's syndrome, a condition that often occurs in patients with bilateral damage to pareto-occipital areas of the cortex.[9] Balint's syndrome is interesting for present purposes because the visual systems of subjects with this condition can seem to be non-photographic examples of spatially agnostic (visual) informants. Patients with Balint's syndrome can have normal visual acuity and full visual field, and can report successfully the color and shapes of seen objects. However, they are afflicted by three main visual pathologies: "simultanagnosia (the inability to see more than one object at a time); optic ataxia (the fixation of gaze with severe problems in voluntarily moving fixation); and optic apraxia (the inability to reach towards the correct location of perceived objects)" ([Robertson, 2003], 96). In addition (presumably because of the optic ataxia associated with the syndrome) such subjects are extremely poor at recognizing and reporting the geometric arrangement of objects in space or estimating distances in both 2D and 3D visual space. These facts might suggest (indeed, have suggested to more than one commentator on our earlier work) that the visual systems of Balint's syndrome patients are spatially agnostic informants; if so, our view entails that these patients fail to see objects. But this seems hard to accept; indeed, it flies in the face of the most natural way to describe the condition (viz., that subjects see only one object at a time and cannot relate to this object spatially).

We are inclined to think that Balint's sufferers do see objects (as the natural description bears out), and are therefore committed to saying that their visual systems carry e-information. But why suppose their visual systems *don't* carry e-information? Not because these subjects are unable to report on the location of objects in egocentric space: as we have emphasized, the ability to report locations in space is not required of subjects whose visual systems carry e-information. Not because they are unable to reach toward objects: nowhere have we suggested that successful performance on reaching tasks is criterial for carrying e-information. Rather, we have said, what is criterial for carrying e-information is an objective, probabilistic, counterfactual-supporting link between two independent events. As far as we can see, the obvious explicit performance measures (verbal report, reaching behavior) fail to show either that

[9]Thanks to Lynn Robertson for advice regarding our discussion of Balint's syndrome.

there is or is not the informational link in Balint's syndrome visual systems that we claim there is. The evidence considered so far leaves open the following two possibilities: (i) Balint's visual systems fail to carry e-information about seen objects, and this is what leads to the difficulties with reporting or reaching toward the locations of seen objects, or (ii) Balint's syndrome visual systems do carry e-information about seen objects, but the subjects fail to report on or reach correctly for seen objects because this e-information cannot be integrated with speech and motor centers. However, a number of recent experiments using implicit measures of spatial information strongly support interpretation (ii). For example, [Kim and Robertson, 2001] found systematic effects on reaction times in spatial alignment tests for both normal and Balint's syndrome subjects. In related work, [Robertson et al., 1997] report that reading time of the words 'up' or 'down' increased if those words were presented in a location at odds with its meaning ('up' presented at the bottom of a rectangle, 'down' presented at the top of a rectangle) relative to reading times when the same words were presented in the opposite locations (locations consonant with the word meanings); similar results are reported by [Humphreys and Riddoch, 2003] (cf. [Robertson, 2004], 164–177). In all these cases, the locations of distal items seems to have systematic (though implicit) effects in Balint's patients, and it is hard to see how this could be true (barring action at a distance) unless some state of (some part of) the visual system carried information about locations of distal items. For this reason, it seems reasonable to think that Balint's syndrome visual systems carry e-information, hence that they satisfy our requirement on object seeing.

Fair enough. But we can surely imagine a condition that is just like Balint's syndrome in its symptoms, but such that the visual systems of its sufferers do not carry e-information. Call this imagined condition 'Shmalint's syndrome'. We are committed to saying that Shamlint's patients fail to see objects, since, by stipulation, they fail to satisfy our necessary condition on object seeing. But, on the other hand, the initial reason given for holding that Balint's patients see objects seems to apply to Shmalint's patients too: the most natural way to describe the condition is to say that patients see only one object at a time but are impaired (to different extents) in relating spatially to it. We are prepared to bite the bullet here and insist that (the naive description notwithstanding) Shmalint's patients fail to see objects. It is worth bearing in mind that apparent Shmalint's patients may turn out to see objects after all: absence of evidence of e-information (e.g., absence of the sort of implicit evidence of e-information we found in Balint's patients) cannot be taken as evidence of absence of e-information. But, we claim, true Shmalint's patients — viz., patients whose visual systems, as a matter of fact, fail to carry e-information — fail to see objects. Of course, the visual systems of Shmalint's patients will carry v-information, e.g., about the color and shape of a single object in front of them. Perhaps this is enough to say that the knowledge derived therefrom counts as perceptual knowledge. Nonetheless, our contention is that, insofar as they fail to carry e-information, what these visual systems are doing falls outside what

counts as object seeing.[10]

3.2 Questions About Information

We have argued (1) that photographs provide v-information, but (2) they do not provide e-information. Nevertheless, we would not be surprised if readers were still a bit suspicious of both of these claims. With respect to (1), it is indisputable that some photographs mislead. For example, some photographs distort the shapes and sizes of the objects they depict. And photographic color is highly dependent on film processing, so processing errors may lead to inaccurate depiction of object color. Hence, one might be tempted to deny that highly distorted and poorly processed photos provide v-information. With respect to (2), some photographs seem to enable us to judge the distance between ourselves and the objects they depict. Some commentators on our earlier paper have suggested that these photographs provide e-information.

We believe that both of these objections are misguided. In both cases, the objections rest largely on a misunderstanding of our claims. The crucial point is that our account of information is utterly non-doxastic. That is, we do not understand the provision of information in terms of a capacity to provide true beliefs nor, in fact, do we construe information in any cognitive terms at all. We lay no claim to the term 'information' however. So it may be that, in some sense or other of 'information', there are photos that do not provide visual information about the objects they depict, and that there other photos that do provide e-information about their depicta. Once the use of term 'information' is disambiguated, the truth of these claims will be clearly shown to be irrelevant to (1) and (2). Let us then turn to consideration of criticisms of our claim that photographs provide v-information.

It is certainly true that some photographs are misleading, insofar as they

[10]The conclusion that Shmalint's patients fail to see objects is, as we have said, at odds with the naive description of the case, and is to that extent a counterintuitive result of our account. We are prepared to live with this counterintuitive result because, after all, Shmalint's syndrome is an extremely abstruse (not to mention fictional) condition.

However, there is another line of response for those with weaker stomachs for counterintuitive results. Namely, one might reformulate our necessary condition on seeing this way: a token process counts as seeing x only if properly functioning tokens of that process carry e-information about x. The thought, then, is that a theory of proper functioning would help us type token processes in a way that would allow a little more flexibility about cases. In particular, the contemplated revision would allow that Shmalint's sufferers see objects so long as their token visual processes turn out to be improperly functioning tokens of the type of visual process in non-Shmalint's subjects (and assuming that tokens of this type of visual process carry e-information). On the other hand, looking at photographs would still fail to carry e-information, hence would still fail to facilitate object seeing. (Presumably, if you prefer this line on Shmalint's syndrome, you would say exactly the same thing about Balint's syndrome as well.)

Of course, the price of this alternative account is the provision of an acceptable theory of proper function. We realize that many philosophers are prepared to pay this price, even though most of them regard it as large. However, given the tiny benefit to be had in the context of the present discussion — a more intuitive description of a fanciful visual pathology, it is unobvious that the price is worth paying for present purposes.

tend to lead viewers to form misguided beliefs about the objects they depict. In particular, many photographs mislead with respect to the visually detectable properties of the objects they depict. A photograph of a house (or horse) may make it look small smaller than it really is. A photograph of your grandfather may make him look more handsome than he really was. And the poorly processed (and poorly lit) photographs of your nature hike may make the scenery look more blue than it was when you were there. But these indisputable truths are irrelevant to the question of whether such photographs provide v-information, given our technical understanding of 'information'. If the horse had been smaller, then the image would have been different. If your grandfather had been more handsome, then the photo would have looked different. So the requisite informational link is still present in these cases. What the cases show is that the process of forming beliefs about the visually detectable properties of objects on the basis of photographic images of them is not perfectly reliable. Providing v-information is not sufficient for producing a true belief. Who would have thought otherwise? Certainly not us.

(There may be some other sorts of photographs that do not provide v-information. For example, if there are non-depictive photographs then such photographs do not provide v-information. But this is not damaging to any view we are propounding. Our aim has been to explain the distinctive epistemic status of photographs qua type of depictive representation. The class of photographs under consideration, then, is properly contained in the class of depictive representations. As such, non-depictive photographs lie outside the range of our explanatory target.)

What about criticisms of (2) — our claim that photographs do not provide e-information? One reason for doubting this claim might stem from the thought that photographs "can serve, along with information from other sources in an inference to egocentric information" ([Currie, 1995], 66). For example, consider a photograph of a well-known landmark (e.g., the Buddy Holly status in Lubbock, Texas). If an agent is confronted with such a photograph, and she knows both her location and the location of the landmark, it may be possible for her to calculate the location of the depicted object with respect to her current location. Nonetheless, we contend that the photograph in question fails to carry e-information. After all, if the agent moves with the photograph, the egocentric location of the landmark changes, but the photographic image does not change. The photograph itself does not carry e-information.[11] (Of course, the agent may be able to recalculate the egocentric location of the depicted object after moving with the photo. But this is irrelevant to the issue at hand.)[12]

A second reason for doubting our contention that photographs do not carry e-information has been suggested by Kendall Walton [personal correspondence],

[11] It should be noted that Currie recognizes that — despite their occasional role in inferences of the aforementioned kind — photographs do not convey e-information ([Currie, 1995], 66).

[12] Compare the fact that you may be able to calculate the egocentric location of someone who is speaking to you on the telephone if you happen to know where that telephone is and where you are in relation to it. Still, the sounds you hear do not carry e-information about the person to whom you are speaking. The reasoning is the same as in the photographic case.

who has offered a range of cases in which it appears that photographs do, or could (under the right conditions), provide e-information. Walton first asks us to consider Polaroid photos that develop, pop out, and disintegrate rapidly. (Note that such photos would not be very useful evidentially since any information they provide comes with significant strings attached!). He suggests that these photographs might plausibly be said to provide the spatial information that the objects they depict are in the vicinity. He also suggests that current photographs, and in particular, some photographs made before the development of the telephoto lens and space travel, look to provide at least one significant piece of e-information — the information that the object depicted in the photograph occurred on the same planet that the viewer of said photograph inhabits.

We contend that Walton's cases are not, in fact, ones in which photographs provide e-information — at least of the sort that we are interested in. Rather, these cases are similar to the ones that Currie describes above — they are cases in which certain photographs (and the information they do carry) may be used inferentially to form true beliefs about the egocentric location of the objects they depict.

Let us consider the disintegrating Polaroids first. While it is true that in ordinary cases one could reliably infer from such a photographic image that the object was nearby, the requisite informational link does not hold. So long as the photograph has any non-infinitesimal life-span, it is not the case that if the object were in a different egocentric location then the image would be different. For no matter how short a time that the photograph persists, it is still possible for the object or photograph to move without a change in the image. So the photograph does not carry e-information.

With respect to Walton's world-bound photograph case, it may be instructive to consider what enables typical viewers of photographs to know that that the objects depicted inhabit the same planet. Clearly what does much of the work here is the mere artifactuality of photographs. Presumably, prior to space travel it was reasonable to infer that any artifact one came across was made on the same planet that one inhabits. So, prior to space travel, it would be reasonable to believe that any photograph that one came across was made on this planet. While we now have photographs and other artifacts that are not from this planet (e.g., they are made in space), they are still relatively rare. Unless we are in a context where extra-planetary photographs are likely to be displayed (e.g., planetary geology texts), we can assume that the photographs that we are confronted with were made on this planet.

Now it does not follow that any object depicted in such a photograph inhabits this planet, not even if the photo was taken prior to the development of the telephoto lens. Earth-bound photographs can certainly represent astronomical phenomena even without telephoto lenses. But perhaps Walton's point should be understood to only apply to photographs of non-astronomical objects: "Well, think of photographs before long telephoto lenses were devised, when it is clear from qualities of the picture whether the object was photographed from a position on the same planet, or not" [Walton].

Nonetheless, these are still not photographs that provide e-information of

any sort. Presumably one might infer from the look of a photograph that it depicts something that is on this planet. But the counterfactuals relevant for evaluating whether an informational link is present still suggest that the photograph is spatially agnostic. Perhaps what makes this case seem to be one in which e-information is carried by the photograph is that there is a de facto correlation between the photograph being viewed on planet P and the objects it depicts being located on planet P. But, as we pointed out above in the case of closed circuit video, de facto correlations do not entail the presence of informational links. If, contrary to fact, one were able to move to another planet, the photographic image would not change.

3.3 Egocentric/Allocentric Revisited

Our explanation of the epistemic status of photographs rests largely on the claim that photographs carry v-information without carrying e-information. But why formulate this point in terms of agnosticism about *egocentric* location (information about the location of the depictum with respect to the viewer) rather than allocentric location (information about the location of the depictum with respect to some frame of reference independent of the viewer)? When we considered this issue before, we decided to stick with the egocentric formulation for the reason that this allowed us to mark the distinction we wanted to make while avoiding further, controversial issues. Here is what we said then:

> . . . 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 non-prosthetic vision on the other. We suppose we could defend an allocentric/absolute formulation of our requirement if we were willing to rule out the individuative standard at issue or plump for a *de re* reading of the counterfactuals, but we'd 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 ([Cohen and Meskin, 2004], 9–10).

We now think that this was the wrong way to approach the issue; in particular, it seems that the counterfactual question we used here as a test for the presence of an informational relation — a test that works well in a wide variety of situations — is inapplicable to cases of this sort. The relevant notion of information, originally from [Shannon, 1948] (although many philosophers learned

about the idea from [Dretske, 1981]), amounts to a reduction in entropy, or uncertainty. As such, if some outcome is certain, there is no entropy about that outcome that can be reduced, so nothing can carry information about that outcome. This is expressed in Shannon's setup by noting that if a random variable p has only a single outcome with non-zero probability, then the entropy for that variable is already zero, so can't be reduced; it follows that the maximum mutual information between p and any other random variable must be 0. Hence, on this setup, nothing can carry information about p. Dretske's (slightly different) setup gives the same result; for him, a signal r carries the information that p just in case the conditional probability of p, given r, (and k, the knowledge of the receiver of r), is 1 (but given k alone, less than 1) ([Dretske, 1981], 65). But if there is no alternative outcome for p, then the conditional probability of p given k alone is unity; hence, on this definition, too, nothing can carry the information that p.[13]

Of course, informational relations support counterfactuals. For this reason, it is usually possible to test for the presence of an informational relation between p and q by asking about this counterfactual: if q were different, p would be different. However, the lesson of the foregoing is that this method will lead to errors in the case where q is necessary. As explained above, if q is necessary then nothing can (*a fortiori*, p cannot) carry information about q; but the standard semantics has it that the counterfactual comes out vacuously true. Consequently, the counterfactual test yields the incorrect verdict that there is an informational relation in such cases, and so cannot be trusted in these cases.[14] (The counterfactual test is, as far as we can see, safe in other cases; hence we continue to appeal to it in this paper.)

But the reason we gave, in the earlier paper, for preferring a condition formulated in terms of egocentric rather than allocentric location turns crucially on using counterfactuals as a test for informational relations in cases where this test will go wrong. We had said that an allocentric formulation is undesirable because it makes for an informational relation (barring controversial stipulations we wanted to avoid) where none is wanted; and we said this because the allocentric formulation resulted in vacuously true counterfactuals. Now that we see that the counterfactual test delivers the wrong verdict in such cases (viz., cases where the counterfactual comes out vacuously true), and that in fact informational relations can't hold in such cases, it is clear that the an allocentric formulation precludes an informational link (again, barring controversial stipulations we want to avoid making), as desired. The upshot is that our reason for

[13]Thanks to Eric Thomson for help on the information theory.

[14]A particularly egregious instance of the same difficulty concerns viewing photographs of oneself. If the photograph's depictum and viewer coincide, then egocentric location of the depictum with respect to the viewer is necessarily fixed at the origin. Consequently, the counterfactual ordinarily used to test for the presence of an informational relation will come out vacuously true. Again, if we took this to show that photographs of x carry e-information when viewed by x, then our account would allow that a photograph of x is transparent when viewed by x (but not when viewed by y when $x \neq y$). We do not accept this consequence. Rather, we take this result as another demonstration of the inapplicability of the counterfactual test for informational relations in cases where the counterfactuals are vacuously true.

preferring an egocentric formulation has evaporated.

So now we can ask again: is there any reason for preferring a condition formulated in terms of information about egocentric location (e-information) rather than information about allocentric location (a-information)?

One reason to prefer a formulation of the epistemic status of photographs in terms of e-information rather than a-information harks back to our goal of distinguishing processes that constitute object seeing from processes that provide v-information without thereby allowing for object seeing. Consider the process of ordinary, non-prosthetic object seeing; while tokens of this process type often carry a-information about seen objects, plausibly there are tokens of the same process type that do not. For example, some seen objects move in unison with the seer (e.g., a ring on one's finger or the eyeglasses perched on one's nose). As I move in allocentric space while attending to my eyeglasses, the eyeglasses change their a-location even though my visual image of the eyeglasses is unchanged. This suggests that some process tokens we want to count as instances of object seeing fail to carry a-information about seen objects. But recall that we hoped to distinguish the process of acquiring v-information via photographs from the process of acquiring v-information via object seeing by appeal to locational information carried by those processes. The present reflections show that, if 'locational information' is understood in terms of a-information, then the process types are alike in the locational information they carry, contrary to our aim. In contrast, if 'locational information' is understood in terms of e-information, then the two process types come apart, as desired.

4 From *Blow-Up* to *Blow Out*: The Epistemic Status of Sound Recordings[15]

We have claimed that the distinctive epistemic status of photography lies in the fact that it is a categorially salient and spatially agnostic source of information. Let us now consider another technology — sound recording — that appears to share these features with photography. Just because it shares these features, our analysis predicts that sound recording should have a similar epistemic status as that of photography. And, as we argue below, this is precisely the case.

Sound recordings are significant sources of information. Of course, they do not typically provide information about the visually detectable properties of the objects they represent.[16] Rather, recordings typically provide information about

[15]Michelangelo Antonioni's 1966 classic, *Blow-Up*, concerns a photographer who discovers that he may have accidentally photographed a murder. Brian DePalma's 1981 thriller, *Blow Out*, inspired by Antonioni's film, concerns the accidental sound recording of what may have been the murder of a presidential candidate. Both films play on the distinctive epistemic status of spatially agnostic informants.

[16]Here we are assuming that audition and sound recordings, like vision and photographs, represent individual objects. An alternative choice would be to say that audition and sound recordings represent events rather than individual objects. (We take it that the same choice is in principle available for vision and photographs, but that it has seemed less tempting for these forms of representation than for audition and sound recordings.) We are officially

the aurally detectable properties of the objects they represent (i.e., they provide what we shall call 's-information'). Evidence for this can be seen in the usual sort of counterfactuals. If the musician had played different notes on the piano (and, hence, the performance had been different), the musical recording would have sounded different. If the speaker had said different words (i.e., if the speech were different), then the recording of the lecture would have sounded different. But sound recordings, like photographs, do not provide e-information. Again, the evidence for this is to be found in the falsity of the relevant counterfactuals. It is not the case that if the Walkman-equipped listener were to change her location with respect to the (recorded) auditory object that she would hear something different.

So sound recordings are spatially agnostic informants. They provide s-information without providing e-information. And this distinguishes them from ordinary hearing — and ordinary auditory prosthetics such as hearing aids — in just the way that photographs are distinguished from ordinary seeing. For ordinary hearing also provides us with e-information. Consider ordinary conversations. As your interlocutor moves with respect to your location, what you hear is different. For example, the closer she is, the louder the sound you hear. In addition, ordinary hearing is such that the provision of s-information is tied to the provision of e-information. If you are not in a position to get e-information about an object through ordinary audition you will typically not be able to get s-information about that object from it either. Sound recording, then, bears a relation to ordinary hearing that is analogous to the relation photographs bear to ordinary seeing. While photographs provides v-information in contexts in which e-information is unavailable, sound recordings provide s-information in contexts when e-information is unavailable.

One apparent difference between photographs and sound recordings has to do with their categorial salience. While photographs are typically categorized as photographs on the basis of purely visual cues, sound recordings do not wear their category on their sleeve, as it were. While we are almost never fooled by a photograph into thinking we are directly looking at what it depicts, sound recordings often confuse us. How many times have we thought we heard a person talking in the next room only to find out that it was merely a recording? Still, in contexts in which the category of audio recordings is salient they are treated as having a distinctive epistemic value.

Thus, we claim that sound recordings are, like photographs, categorially salient and spatially agnostic informants; therefore, our theory predicts that sound recordings will have a special epistemic status. And, in fact, this is the case. Just as photographs serve as both formal and informal evidence, so too do sound recordings. In a legal context, undercover agent wears a wire in order to acquire evidence about criminal misdeeds. Phone conversations are tapped for a similar purpose. Old recordings of tribal musical performances provide evidence for ethnomusicologists. And recordings of our loved ones provide informal

agnostic about this issue. Therefore, if you prefer the second way of talking, feel free to regard our talk about audition representing an object x as a shorthand form for talk about audition representing the event of x's exemplifying some relevant property.

evidence of what they sounded like.

While one might be tempted to explain the distinctive epistemic status of sound recording by reference to some auditory version of the transparency thesis (i.e., the view that sound recordings allow us to literally hear the objects recorded), we believe that there is no need to plump for such a view. It is sound recording's status as a categorially salient and spatially agnostic informant that explains its distinctive evidentiary role.

References

[Bazin, 1967] Bazin, A. (1967). *What Is Cinema?* University of California Press, Berkeley. Translated by Hugh Gray.

[Cohen and Meskin, 2004] Cohen, J. and Meskin, A. (2004). On the epistemic value of photographs. *Journal of Aesthetics and Art Criticism*, 62(2):197–210.

[Currie, 1995] Currie, G. (1995). *Image and Mind: Film, Philosophy and Cognitive Science*. Cambridge University Press, Cambridge.

[Dretske, 1981] Dretske, F. I. (1981). *Knowledge and the Flow of Information*. MIT Press, Cambridge, Massachusetts.

[Goodman, 1976] Goodman, N. (1976). *Languages of Art: An Approach to a Theory of Symbols*. Hackett, Indianapolis.

[Humphreys and Riddoch, 2003] Humphreys, G. W. and Riddoch, M. J. (2003). From *What* to *Where*: Neuropsychological evidence for implicit interactions between object- and space-based attention. *Psychological Science*, 14(5):487–492.

[Kim and Robertson, 2001] Kim, M.-S. and Robertson, L. C. (2001). Implicit representations of space after bilateral parietal lobe damage. *Journal of Cognitive Neuroscience*, 13(8):1080–1087.

[Robertson, 2003] Robertson, L. C. (2003). Binding, spatial attention, and perceptual awareness. *Nature Reviews Neuroscience*, 4:93–102.

[Robertson, 2004] Robertson, L. C. (2004). *Space, Objects, Minds and Brains*. Psychology Press, New York.

[Robertson et al., 1997] Robertson, L. C., Treisman, A. M., Friedman-Hill, S. R., and Grabowecky, M. (1997). The interaction of spatial and object pathways: Evidence from Balint's syndrome. *Journal of Cognitive Neuroscience*, 9:691–700.

[Shannon, 1948] Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656.

[Walton, 1984] Walton, K. (1984). Transparent pictures: On the nature of photographic realism. *Critical Inquiry*, 11:246–276.