

Justification as the appearance of knowledge

Steven L. Reynolds

© Springer Science+Business Media B.V. 2011

Abstract Adequate epistemic justification is best conceived as the appearance, over time, of knowledge to the subject. ‘Appearance’ is intended literally, not as a synonym for belief. It is argued through consideration of examples that this account gets the extension of ‘adequately justified belief’ at least roughly correct. A more theoretical reason is then offered to regard justification as the appearance of knowledge: If we have a knowledge norm for assertion, we do our best to comply with this norm when we express as assertions only beliefs that appear to us to be knowledge. If we are doing our best, there is little point in further sanctions. So a norm of knowledge for assertion would lead to a secondary norm of justified belief as the appearance of knowledge, marking a point at which our assertions may be corrected but should not be blamed.

Keywords Knowledge · Justification · Epistemic norms · Appearance

The thesis of this paper is that adequate epistemic justification is the appearance of knowledge.¹ I believe with justification that I am currently working on this paper if and only if there has been an appearance to me of my knowing that I am currently working on this paper.

This is offered as an account of doxastic justification, not of propositional justification. It holds that a belief is adequately justified for S if and only if, over the

¹ Previously proposed, but not defended, in Reynolds (1998, pp. 533–534, 2002, pp. 151–152).

S. L. Reynolds (✉)
School of Historical, Philosophical, and Religious Studies, Arizona State University,
P.O. Box 874302, Tempe, AZ 85287-4302, USA
e-mail: Steven.reynolds@asu.edu

time in which S acquired and maintained it, it appears to S to be knowledge. It does not say that a belief that p is justified for S iff it would appear to S to be knowledge that p were S to acquire or sustain it by certain readily available cognitive processes, as in ‘propositional’ justification. It is also an account of adequate or full justification, not of justification as a matter of degree. There aren’t degrees of appearing to be knowledge.

There are some rival views of adequate justification at about the same level of generality. A deontological view of the nature of adequate justification claims that justification is believing as one ought to believe, all (epistemic) things considered (BonJour 1985; Chisholm 1991; Kornblith 1983). There is some question whether ‘ought’s and ‘should’s as applied to beliefs indicate epistemic duty (Alston 1985, 1988; Plantinga 1993a). Perhaps they indicate instead proper functioning (Plantinga 1993a, b) or acceptable performance of a role (Conee and Feldman 2001; Feldman 2001, 2004, 2008). A second view of adequate justification is that it is believing in a way that appropriately advances our cognitive goal of true belief (Alston 1985; BonJour 1985, 1998). On this view adequate justification is usually said to require an internally available, and on some views internally defensible, indication that our belief is likely to be true, a criterion of truth. A third view suggests that we are justified in a belief if and only if our method or process of arriving at the belief is such as to reliably yield true beliefs in our normal environment (Goldman 1979). No internally available criterion of truth is required on such reliabilist views. The view that justification is the appearance of knowledge may be compatible with these views, if they are regarded merely as generalizations about justification, but it is a rival to them as an account of the nature of adequate justification. I hold that adequate justification cannot be usefully analyzed or explained in terms of a goal of true belief, or responsible inquiry on the part of the believer, or reliability in the believer’s processes of coming to believe.

On the appearance of knowledge view, justification is to be understood in terms of knowledge. I agree with Williamson in thinking that knowledge is a basic concept, not to be given a definitional analysis (Williamson 2000). The concept of knowledge does not have epistemic justification as a definitional component, although justification is a necessary condition for knowledge. We learn to judge that someone knows by observing and considering examples of knowledge and non-knowledge and hearing criticisms of others’ judgments whether someone knows, not by hearing characterizations of knowing that relate it definitionally to notions of truth, belief, evidence, or justification (etc.). We come to be able to recognize knowledge, in something like the way we learn to recognize faces, or kinds of birds, or that someone is anxious about a test, without being able to say much about how we do it. This is compatible with holding that we can, on reflection or after other inquiry, say a lot of interesting and substantive things about knowledge and its relation to evidence, belief, perception, testimony, memory and the like, just as we can say quite a lot about human faces’ relations to their noses, cheeks, eyes, lips, foreheads, and the underlying bone/cartilage structures.

1 The appearance of knowledge

The key term for the view that justification is the appearance of knowledge is ‘appearance’.² It’s one of the more problematic philosophical terms, and the way I want to use it is inconsistent with some recent ways it has been used, especially in the epistemological literature.³ There the focus has recently been on transitions from an appearance that *p* to a belief that *p*, and it is thought that one can’t have an appearance that *p* unless one has the concepts to believe that *p*.

But I have in mind a perfectly ordinary and even common use of the term ‘appearance’ which doesn’t require that the subject have such concepts. Consider the following example: Walking in a park I notice an unfamiliar bird, and decide I would like to find out what it is. Fortunately it doesn’t immediately fly away, so I observe it for two or three minutes. A few hours later, having returned home, I look up a web site, find a few photos, follow up by watching a video, and conclude confidently that I saw a Steller’s Jay. I think it is perfectly correct to say that the bird I saw had the appearance of a Steller’s Jay, even though I didn’t know that that’s what it was at the time. If it hadn’t had the appearance of a Steller’s Jay, I wouldn’t have been able to remember that appearance later and match it to the photos and video of Steller’s Jays. I didn’t have the concept of a Steller’s Jay, yet I had an appearance of a Steller’s Jay.

Another example: I become familiar with the appearances of my students during class, then look them up by picture on the class roster and so learn to identify them by name. It would be misleading to say that it appeared to me that Sandra asked a very astute question earlier today. That would suggest that I identified her as Sandra when she asked the question. But I did have an appearance of Sandra asking an astute question, or I wouldn’t have been able to learn her name by looking at the photos.

It is said that those of us who are red-green color blind lack the standard concepts of green and red. Imagine that I am given gene therapy to correct my red/green deficiency, as has recently been done for some adult squirrel monkeys (Mancuso et al. 2009). My eyes are bandaged (for the sake of philosophy) from the beginning of the therapy until the doctors are confident that my eyes now have the formerly missing variety of cones. The bandages are removed and I see an unlabeled large square sample of a color quite different from any I have seen before. I don’t know whether I am seeing red or green, since it looks quite different from the previous samples of those colors as I have seen them. After enjoying the new color sensation for a few minutes, I walk into the next room and look at some labeled color samples,

² Smithies (2011) suggests that justified belief is subjectively indiscriminable from knowledge, which sounds close to what I am suggesting, but he endorses a cognitive account of indiscriminability: “the subject in question cannot know that he instantiates the one state while also knowing that he does not instantiate the other”. The claim that justified belief is subjectively indiscriminable from knowledge is apparently treated as a substantive assertion about the relation of the two, rather than an attempt to give a reductive account of epistemic justification.

³ On some views beliefs that, as it were, transcribe the content of an associated appearance are said to be *prima facie* justified, regardless of the type of appearance involved (Huemer 2007) or at least *prima facie* justified if the appearance is a perceptual appearance (McDowell 1994; Pryor 2000). That is not the sort of appearance I have in mind.

pick out the color I have just seen, and discover that, when I first opened my eyes I then had for the first time an appearance of green. It seems that I didn't acquire the standard concept of green until I looked at the labeled color samples, since when the bandages were first removed I still couldn't tell, on the basis of vision alone, that I was seeing green. Yet I must have had an appearance of green during the first moments after the bandages were removed, or I wouldn't have been able to recognize that appearance when I subsequently saw it again on a labeled color sample. So having an ordinary appearance of green does not require having the standard concept of green. The sample looked green, to me, when the bandages were first removed, even though I lacked the relevant concept of green.

Similarly, I hold, small children, or mentally disabled people, or even animals, who don't have the concept of knowledge, may nevertheless have an appearance of themselves knowing. My dog sees me picking up the dog dish. I think she also then has an appearance of herself knowing that I am picking up the dog dish. How can that be, given that she doesn't, and probably can't, have the concept of knowledge? Having the concept isn't necessary for having the appearance, as we have just seen in these examples. She has the visual appearance of me picking up the dish, and she doesn't have any awareness of undermining beliefs or a lack of confidence that I am picking up the dish (and presumably couldn't, since dogs can't have those sorts of self-awareness). When I'm seeing someone pick up a dish, and lack the sorts of undermining awarenesses that the dog also lacks, I have an appearance of knowing someone is picking up the dish. So she also has what we should count as an appearance of knowledge. But since I also have the concept of knowledge and some capacity to identify instances of knowledge, I can come to know that I know, which of course the dog cannot do.

So 'appearance' as I shall use the term is not meant as a synonym for 'believes' or 'thinks' or even 'seems'. It couldn't be, as appearances in the relevant sense may be had by those who lack the concepts required to have the corresponding beliefs or thoughts.

It is quite possible to be mistaken about appearances, although that would be rare if the appearance were very simple, as in the color example. In the imagined case I might have wrongly thought I was seeing red. It would be an ordinary matter for me to think I was having an appearance of a Blue Jay when I wasn't, but was instead having an appearance of a Steller's Jay. The correction in that case wouldn't have to be based on a second look at the bird. I might correct myself by just remembering what the bird had looked like, wondering perhaps how I could have been so far mistaken at the time. There may be appearances that are infallible or incorrigible, as philosophers have said, but they're not the sorts of appearances I am using to characterize epistemic justification.

In the hope of averting another sort of mistake about my project: I have been arguing that one can have an appearance of X F-ing without have the requisite concepts to believe that X F-s. I have not of course argued that the appearances I do have in those cases involve *no* concepts. My observations of the bird that produced my ability to recognize photographs and videos of that type of bird later invoked concepts such as 'bird' and 'color' and so on, which I used in thinking about what I was doing then and later. Even my imagined first experience of the color green

involves thinking of what I then saw as a color. Perhaps I also used the concept ‘red-or-green’ (the color concept we red-green deficient persons use), when I thought “That’s not really very much like red-or-green, is it?” I mention this point mostly to ward off impressions that I think I’ve successfully argued for something I haven’t even tried to argue for here, such as conceptual-content-free experiences.

What can we say to further clarify the notion of an appearance of knowledge? It is an important motivation for the view that we find that we can’t give a descriptively adequate account of how we can tell that someone knows. We are typically in something like the position we are in with respect to human faces, which most of us are quite unable to describe adequately but are nonetheless very good at recognizing. An appearance of knowledge is typically an awareness of many things—beliefs, feelings of doubt, perceptual experiences, memories, logical and evidential relationships among beliefs, and so on. It is often extended over a significant period of time—minutes or hours or even longer. When someone watches me working on this paper (seeing me typing, reading, pacing, looking at books or journals), there is an appearance to them of my knowing that I am working on this paper. In the same case, in virtue of my experiences while working on this paper, including seeing and feeling what I am doing (typing, reading, handling books, talking to myself, etc.), and my awareness of my deliberations, deciding, and intentions, there is an appearance to me of my knowing that I am working on this paper.

Other typical examples: when I am aware of myself seeing that *p*, there is an appearance to me of my knowing that *p*. When I am aware of myself adequately considering and understanding a proof that *p*, there is also an appearance to me, over the time in which I am considering the proof, of my knowing that *p*. Likewise for cases of receiving testimony and thereby coming to know something.

One reason that might be given for doubting that there are appearances of oneself knowing is that there are views that deny that self-attributions of some psychological states, such as belief, are typically based on self-observation. To the extent that self-attribution of knowledge depends on self-attribution of such states, it seems that the self-attribution of knowledge must also be non-observational. On a certain theory about the requirements for genuine observation that may be true—perhaps the self cannot be an object of observation, but is necessarily only a subject who has observations. Such views are encouraged by puzzling features of self-attribution of psychological states and events and even of self-attribution of some physical states and events such as postures and bodily movements.

Since I have defended an experiential view of these self-attributions elsewhere (Reynolds 1992), I will not review this problem here.⁴ Our judgments that we know, e.g., that there is a cup here, or that this theorem is thus proven, would usually not be made, or would be obviously wrong, in the absence of perceptual experiences of seeing, touching, smelling, etc., or of cognitive experiences of various kinds of reasoning, remembering, etc. Our judgments that we know are usually based on

⁴ See Bar-On (2004), for discussion of such views, and endorsement of a well worked out “avowals” view of self-knowledge.

such experiences, in some good sense of ‘based on’. Even where the knowledge is for example knowledge of our own beliefs, wishes, hopes, and so on, where it may be doubted that there is a distinctive kind of quasi-perceptual experience of those states, there is a kind of awareness of these states, which can be mistaken, and which can be described in a broad sense of the word as an appearance of those states. I propose to regard all of these kinds of subjective and/or experiential access to states or events that may be pertinent to our judgments whether we know as parts or aspects of an appearance of knowledge. Sometimes (I claim) they are collectively an appearance of my knowing that *p*, and sometimes they are enough like such an appearance that their occurrence can explain my mistakenly thinking I know that *p*. In this paper I don’t want to enter any further into the dispute whether this is a proper use of the term ‘appearance’, so I invite the reader who is troubled by doubts on this point to treat this as a stipulation regarding my intended use of ‘appearance’ (although I think it is not only that).

Sometimes we can tell that we know, and sometimes we think that we know although we shouldn’t think so. On the other hand sometimes we are mistaken in thinking we know, but it’s not our fault (as in the Gettier cases—see below). I take these to be plain facts, not controversial philosophical claims (however incompatible with some kinds of skepticism). They are naturally explained by saying that we have appearances of knowledge, appearances that can be misinterpreted or that can be misleading. The reader who doubts that we have such appearances should attempt to explain these facts on her preferred alternative view of how we (sometimes) know that we know.

2 The same extension

I shall now argue that the view that justified belief is the appearance of knowledge gets the extension of ‘justified belief’ at least roughly correct. First, I shall argue that if the subject has an appearance of herself having knowledge that *p*, then she is justified in believing that *p*.

Sarah is driving slowly through the countryside, and she sees a barn about fifty yards off the road. She sees it clearly, it is a barn, and she is very well able to recognize barns. It appears to Sarah that she knows it is a barn.

Now amend the case so that, although the structure Sarah is seeing is in fact a barn, there are many fake barns in the area that are difficult to tell from the real thing if observed from the road (Carl Ginet’s well known example, reported in Goldman 1976). Although most philosophers would deny that in the amended case Sarah knows that the structure she is seeing is a barn, they would still hold that she is justified in believing that it is a barn. The only difference in the amended case is the presence of fake barns in the area. Sarah has never seen or heard about the fake barns, and so does not believe or have any reason to believe that there are any. So it must still appear to her, in this amended case, that she knows that that is a barn. It is not knowledge; but it appears to her to be knowledge, and it is a justified belief.

In a second amendment to this case, everything appears just the same to Sarah, and all of her memories and thoughts and beliefs are the same (or as close as they

can be subjectively, if not quite the same content), but she is looking at a fake barn. Her belief that it is a barn, which has the appearance of knowledge to her, is a case of justified false belief.

A second case: Smith and Jones have applied for a job. Smith has strong evidence that Jones is the man who will get the job and that Jones has 10 coins in his pocket. Smith counted the coins that Jones now has in his pocket a few minutes ago, and he received the personal assurance of the company president that Jones would be selected. Smith believes on the basis of this evidence that the man who will get the job has 10 coins in his pocket. If that is all we know about the case, it appears to us and to Smith that he knows that the man who will get the job has 10 coins in his pocket.

Now let us amend the case. Suppose Smith will get the job, instead of Jones, and he happens to have 10 coins in his pocket, but he doesn't believe either of these facts (Gettier 1963). Evidently he does not (in the story as amended) know that the man who will get the job has 10 coins in his pocket. But from his point of view there is no difference from the previous case. In the case without the amendment his belief has the appearance of knowledge. In the amended case it still appears to Smith that he knows that the man who will get the job has 10 coins in his pocket. Most philosophers agree that he is justified in believing that the man who will get the job has 10 coins in his pocket.

A second amendment would have both Jones and Smith not getting the job, and someone else who lacks 10 coins getting it instead, but everything appearing just the same to Smith up to this point. This would be another example of a justified false belief.

The recipe for constructing such trios of cases is: take a case where S knows that p and enjoys the appearance of herself knowing that p. If the case is amended so that the proposition is no longer known, either because it is not true or because its truth is no longer properly connected to the appearance of knowledge, but S still enjoys, in the same way, the appearance of knowledge that p, then S will be justified in her belief that p. I conclude that if there is an appearance to S of herself knowing that p, then she is justified in believing that p.

There are cases where the (real) appearance of knowledge guarantees the fact of knowledge, so that we can't construct the second and third, non-knowledge, cases. For example, the appearance may involve really following a valid mathematical proof from previously known premises. (Merely thinking one is following a valid proof, where the proof is in fact invalid, is not a case of it appearing to one that one is following a valid proof. This is one of those cases where one can be mistaken about an appearance.) In such cases there is no possible case of non-knowledge that really appears, in the same way, to the subject to be knowledge. But if every possible instance of knowledge is also a justified belief, as is commonly held (justification is necessary for knowledge), then it seems that these cases will not be problematic for my thesis that the appearance of knowledge is the justification of the corresponding belief. For such cases will necessarily also be cases of justified belief.

Now I will argue that if S is adequately justified in believing that p, then she will have had an appearance of herself knowing that p. It will be convenient to argue for this thesis in the contrapositive form, considering cases in which S's belief falls

short of appearing to her to be knowledge and then consulting our intuitions whether she is or is not adequately justified in believing. We will find that, in a broad range of cases, if there is not an appearance of herself knowing, then she is not adequately justified in believing. The survey will indicate (but not demonstrate, since it is just a survey) that the claim is true in every case.

Consider the barn case again. If Sarah became aware of anything that made it not an appearance of herself knowing it was a barn, that difference in appearance would also make her unjustified in believing that it was a barn. Thus if the structure she was looking at looked to her too flimsy to be a real barn, or if she thinks she may have heard that people have been building elaborate fake barns in the area, then she would no longer have an appearance of herself knowing it was a barn. She would also not be justified in believing that it was a barn.

Most cases where the subject lacks an appearance of herself knowing are very obviously also not cases of being justified in believing. If S guesses that p, or has made some mistake in reasoning to the conclusion that p, or has been guilty of some noticeable carelessness in observing whether p, or in receiving testimony that p, then there will not be an appearance for S of herself knowing that p. (Reminder: she may of course believe or think that she knows, in such cases, in spite of not having an appearance of herself knowing. We often make mistakes about more complex appearances.) It is clear in such cases that she is not justified in believing that p.

But there are some cases of failing to have the appearance of oneself knowing that are not so obviously cases of lacking justification. I have forgotten the sources of many of my beliefs about current events, geography and standard history, even though those beliefs were acquired from good newspapers, teachers, text books, or other respectable sources. Other beliefs I currently hold have no such respectable source, but I've forgotten this. So these beliefs all equally appear to me, *at the moment when I'm attempting to evaluate them*, to be knowledge, or not to be knowledge.⁵ Yet it seems that we count many standard beliefs about history or geography (etc.) as known, and therefore as justified, in virtue of their actual history, in spite of the absence of an appearance at present that we know them.

In order to rule properly on these cases, my account requires a diachronic, rather than synchronic or momentary, appearance of knowing. Cases where errors in reasoning or perception or receiving testimony have now been forgotten do not count as cases where the belief appears to the subject, over time, to be knowledge, because their actual subjective history will include those errors. The only cases that appear to the subject, over the history of the belief, to be knowledge are those in which the actual subjective histories of the belief instantiate the pattern of knowledge. It seems that even the strictest internalist should want to allow that it is the diachronic appearance of knowledge that is relevant to justification, because only so can she allow that following a lengthy proof justifies the belief in the conclusion when the earlier stages of the proof have been forgotten. Another reason why we would allow a diachronic appearance as of knowledge to count as

⁵ Conee and Feldman argue that we have more confidence in the justified beliefs and that they seem to us more likely to have been supported by a good source. But this seems awfully thin for justification (Conee and Feldman 2001, pp. 69–72).

justification, in the absence of a current appearance of knowledge, will be discussed in the concluding section of this paper.

Here is another example of a belief that does not appear to the subject to be knowledge, yet which many would hold is justified. Suppose Sam has a ticket in a lottery where the odds against his ticket's winning are a hundred thousand to one. It does not appear to him that he knows his ticket will lose. (At least that is the view of many philosophers, who deny knowledge in this case.) But he has very strong reason to believe that it will lose—he knows the odds, and they are enormous. It seems that he is justified in believing that his ticket will lose, but there is no appearance to him of himself knowing it will lose.

One response to this problem case would be to say that, although Sam has strong reason to think his ticket will lose, he is not *fully* justified in believing that it will lose.⁶ At any rate he could have a yet stronger justification, if he were instead aware that the odds were a million to one against and even better justification if there were no chance at all. So in the lottery cases where there is a failure to appear to be knowledge there is also a lack of full or adequate justification.⁷

A more radical response, which I like better but cannot defend here, is to regard probability as part of a different epistemic conceptual scheme, one that does not readily combine with the (older?) scheme that evaluates beliefs in terms of knowledge, justification, and reasons to doubt.⁸ If we can't avoid considering probability of error, as in lottery cases, we combine that consideration with reflection on the question whether there is knowledge by regarding any non-zero probability of error as a reason to doubt, and therefore as undermining knowledge. But in order to retain some role for knowledge evaluations, if the (small) probability of error is not salient we ignore it and attribute knowledge.

Another apparent counterexample: If it is suggested that I may be the victim of a deceiving demon or of an alien super-scientist who now feeds my nervous system radically misleading sensory input, it may appear that, in the absence of a good reason to rule out these apparent possibilities, I lack knowledge, e.g., that here is one hand. Yet although I lack knowledge, I still have strong reason to believe that here is a hand. The Cartesian skeptical hypotheses seem to deprive us of knowledge, but not of justification.

However, the Cartesian skeptical hypotheses are “reasons to doubt”. To doubt is to fail to believe, or at least to believe in a lesser degree (perhaps only very slightly less, where the reason to doubt is, in Descartes' famous phrase, “slight and metaphysical”). So it seems that the Cartesian hypotheses are, in this respect, rather like the lottery cases: they make it appear to the subject that she lacks knowledge, but only by also making the justification for her belief seem less than full or adequate. To the extent that they are reasons to doubt, they are reasons not to believe, and therefore undermine justification for the belief. So this is another case

⁶ See Smithies's distinction between high confidence and belief (Smithies 2011).

⁷ This view of justification thus allows that the knowledge version of the lottery paradox and the justification version should have parallel answers (Nelkin 2000).

⁸ David Lewis expresses a similar view, as have others (Lewis 1999, pp. 439–440).

in which a belief's appearing to the subject not to be knowledge is also its lacking full or adequate justification.⁹

So 'appearance of knowledge to the subject' is at least approximately co-extensive with 'full or adequate justification'. Now I would like to consider a more theoretical reason to regard justification as the appearance of knowledge.

3 Appearance and norms of knowledge

It may seem that the view that justification is the appearance of knowledge offers little hope of contributing to a theoretical understanding of knowledge, since it uses the concept of knowledge, and so cannot be part of an analysis of that concept. I think however that it may help us to understand the functions or purposes of classifying people as knowing. There is a promising candidate for one main function of attributing knowledge, in norms that require knowledge. These are said to include norms for assertion (DeRose 1996, 2002; Hawthorne 2004; Stanley 2005; Unger 1975; Williamson 2000), for practical reasoning (Hawthorne 2004; Hawthorne and Stanley 2008) and for belief itself (Sutton 2007, p. 19; Smithies 2011). Classifying our own beliefs as known, or not known, will enable us to comply with these norms, and classifying others' beliefs will enable us to decide whether they are complying.

According to the thesis that knowledge is the norm of assertion, we are permitted to assert that *p* only if we know that *p*. Let us set aside the question whether we do have such a norm for a moment and consider instead whether we should have such a norm.¹⁰ To answer this question presumably one investigates whether compliance with the norm would be beneficial, whether we can conveniently comply with it sufficiently to achieve those benefits, whether the costs of compliance outweigh the benefits, and so on.

Norms always have costs, at least the cognitive costs of figuring out what we should do to comply and of discerning compliance and failures to comply in others, and the emotional costs of approving and disapproving accordingly. We should have norms only if they carry significant benefits to compensate us for those costs. We will have a norm of assertion only if it seems to us that it tends to produce better assertions, at least on the average and in the long run. Likewise we will have a knowledge norm for practical reasoning only if the cost of having such a norm is compensated, in the long run, by better practical reasoning.

Typically we economize the costs of enforcement of social norms by making a distinction between violations committed deliberately or carelessly and violations the violator had no reasonable way of avoiding. The former sorts of offense reflect unwillingness or insufficient commitment to comply with the norm and so tend to

⁹ Please note that this comment is very limited in its aims. It seeks merely to defuse Cartesian skeptical hypotheses as an apparent counterexample to my thesis, not to suggest, let alone argue for, any particular response to Cartesian skepticism itself.

¹⁰ For those who agree with Williamson's account of what an assertion is, that is a speech act governed by the knowledge norm, this is the question whether we should make assertions or perhaps should instead share information through some other speech act, not governed by the knowledge norm.

draw more stringent enforcement. Thus if a society has a norm requiring respect for the elderly, it matters in that society whether someone respects those who are apparently elderly, whether or not they are in fact elderly, because it indicates whether she is trying to comply with that norm.¹¹ On the other hand, failure to show respect for someone who is in fact old enough to be entitled to it, but doesn't look it, may be excused or mitigated.

So if it is a norm for us that we should assert that *p* only when we know that *p*, it seems that we would show respect for the norm, and indicate to others that we are trying to follow it, by making assertions only when it really appears to us as if we know. If we do that reliably, then our behavior is not likely to be improved any further by blaming or otherwise expressing disapproval of it. If it really does appear to us, in the diachronic sense of 'appear' that I have been describing, as if we know what we assert, but the appearance is misleading and we don't know it, we will still have done our best to comply with the knowledge norm for assertion.

I think that what would be relevant to the assessment whether the subject is adequately trying to comply with the norm of knowledge is the appearance of knowing, rather than, say, the belief that one knows. For we often believe that we know that *p* when we shouldn't.¹² In those cases our assertion that *p* shouldn't be approved by others as an expression of justified belief that *p*. For example, we might have arrived at a belief that we know by endorsing a fallacious inference that led to our belief that *p*, or by trusting a hasty perceptual impression that *p*. In such cases, others who know how we acquired our belief that *p* should not excuse our assertion that *p*, but should instead criticize or at least think the worse of us for making it, even if we happen to believe we know that *p*. So merely believing that we know is not adequate to indicate that we are really trying to follow the norm. On the other hand if we really had an appearance of knowing that *p*, then we could not have been guilty of the fallacious inference or of relying on the careless perceptual impression,

¹¹ A curious case of respect for age: "About 1910 (aged 25) I passed, near Coton in Cambridgeshire, a pair of girls about 8 or 9. One had the face of an angel. Before they were out of ear-shot, angel-face said: I was going to say "bugger" when I saw the old man" (Littlewood 1986, p. 162). Also interesting is Littlewood's reaction on seeing a kindred norm of respect apparently being violated by Hardy, who was then very youthful looking: "In my first year, Dons lunched at the same table as undergraduates, and I once innocently happened to sit next to a block of them. Presently I heard what was apparently an undergraduate chaffing the infinitely venerable Henry Jackson, with great elegance and verve on both sides" (Littlewood 1986, p. 120). We would think less of Littlewood's respect for the relevant social norms if he didn't find this incident remarkable, precisely because we judge whether people are reacting appropriately to the way things appear, not merely to the way they are, and it appeared to him that an undergraduate was engaged in insufficiently respectful conversation with an elderly Don.

¹² Williamson allows that we may reasonably believe that we know, without knowing, and holds that it is reasonable to assert in those circumstances, though still in violation of the norm if we don't really know (Williamson 2000, p. 256). If "reasonable to believe" is understood in the way I recommend for 'justified belief', that is, it is reasonable for *S* to believe *p* if and only if it appears to *S* to be knowledge that *p*, then this suggestion is equivalent to my proposal. Williamson suggests however that what it means is "highly probable on the evidence" and here I am avoiding talk of evidence and the corresponding notion of probability. In discussing a different case he comments that "it is quite reasonable for me to believe not just that there is snow outside but that I know that there is; for me, it is to all appearances a banal case of perceptual knowledge" (Williamson 2000, p. 257). His use of "appearances" rather suggests my own view, but it seems likely to be a matter of speaking loosely, since he evidently thinks the probabilistic inferential relation to evidence is the correct literal characterization.

since such errors would have been part of the appearance to us over time of our acquisition and maintenance of the belief that *p*. So appearance to the subject of herself knowing is more appropriate than belief that she knows for this sort of evaluation.

Do we have a norm of knowledge for assertion? I take this to be a factual question, an anthropological question about our society's practices, or perhaps the practices of all modern societies. Such questions are answered by investigating whether typical or well-socialized members¹³ of the group recognize what the alleged norm requires, whether they comply or attempt to comply with it, whether they criticize, or at least think the worse of, people who don't comply, and are expected to think worse of those who don't comply, and so on.¹⁴

I won't attempt a thorough review of the extensive and growing discussion, but here are two points in favor. First it seems in order to ask someone who asserts that *p* how she knows that *p*. That presupposes that in asserting she represents herself as knowing.¹⁵ Second, it explains why Moore paradoxical sentences of the form “*p*, but I don't know that *p*” are anomalous, although often true. The speaker asserts that *p* but denies that she satisfies the norm for asserting that *p*.

¹³ Of course philosophical theorizing on these very topics may make philosophers atypical, as inclined to give variant answers to protect our theories, or to be just a little uncertain about the customs from being pushed about by the winds of argument, or even just by others' confident claims. If so, we may become unreliable as sources of information about our own customs. But I shall follow the usual (philosophical) custom and hesitantly assume that we are entitled to rely on our carefully considered intuitions, as at least normally an independent check on philosophical theory.

¹⁴ It has been suggested that the evidence whether we have such norms would be “linguistic” evidence (Stanley 2005; Douven 2006). This seems much too narrow, unless one is considering, not whether we have such a norm, but merely whether such a norm is essential to the speech act of assertion. Whether we have a norm of knowledge for assertion would be discovered by considering observations and intuitions of what we do, or would, or should, say, and feel about so saying, in a wide variety of perceptual and cognitive situations. This evidence seems to be much more various and less obviously linguistic than, say, intuitions about sentences' being grammatical, or utterances being phonemically acceptable, or about whether a described speech act counts as an assertion.

¹⁵ Douven (2006, pp. 473–476) suggests that we may so speak because most assertions that satisfy the “rational to believe” norm that (he thinks) really applies will also be such that we know them and it might be less persuasive to note that we don't actually know them. We do sometimes say, in questioning an assertion, “Why do you think so?”. That presupposes that it is rational to think so, rather than that one knows it. But on the knowledge view it makes sense to ask about a necessary condition of knowledge, if we mean to cast doubt on the assertion. If knowledge is the norm of assertion then by asserting one represents oneself as knowing, and therefore also represents oneself as reasonably believing. But if rational belief is the norm, one doesn't represent oneself as knowing by asserting, so it is not so natural to ask how one knows. Compare “why are you so certain?” A mere assertion wouldn't indicate any great degree of certainty, so it would be quite consistent to reply that one wasn't certain, unless something about one's tone of voice in making the assertion suggested certainty. The similar explanation of the Moore paradoxical sentences seems even less plausible, since on the suggestion that we are speaking loosely, taking ‘knows’ as a rough stand-in for ‘rationally believes,’ the Moore paradoxical sentence should merely seem unusual, not paradoxical. Douven acknowledges this, but sets it aside with some comments about how we often find it difficult to make a distinction between pragmatic infelicities such as he claims and inconsistencies such as the knowledge account claims. However, Douven's suggestion that maybe we're getting it wrong in our impression of what's going on, although not obviously false, isn't very convincing as an account of our systematic reaction to these cases. Other things equal, regularity in correct responses seems more likely than regularity in mistakes.

A point against the claim that knowledge is the norm of assertion is that in discussions on topics philosophical, political, religious, sporting (etc.) we often allow assertions that evidently don't express knowledge to pass without criticism, spoken or felt. Thus: "The Lakers are going to win tomorrow" said by someone who has only a fan's knowledge of the team's recent record to go on. Williamson suggests that we have a practice of treating these cases the way we treat grammatical errors in rapid conversation, where corrections are neglected because they are unimportant in the context and contrary to the purposes of conversation (Williamson 2000, pp. 258–259). But there is often no haste or carelessness involved in our failures to restrict our assertions in discussion to what we know, and we seem to have little inclination to withdraw the assertions even on being challenged on that very point. "You don't know the Lakers are going to win." "They will though." Perhaps there is a tacit understanding that assertions in such discussions are to be understood as if prefaced by "I think that...". Or holders of the view might restrict the norm of knowledge to a category of "serious" or "flat out" assertions, perhaps where it is expected that one's word will be accepted.

An argument against the view that knowledge is the norm of assertion has been offered by Jennifer Lackey, as commentary on a series of alleged counterexamples (Lackey 2007):

Suppose that Martin has been raised so that he has strong racist feelings, but now he rejects those feelings and is ashamed of them. Serving as a juror in the case of an African-American man accused of rape, he votes to acquit based on strong evidence that the accused is innocent. But because of his residual racist feelings, he can't quite bring himself to believe in the man's innocence. After the trial is finished, when a childhood friend asks if the man is guilty, he asserts, "No, the guy did not rape her." He doesn't know this, because he doesn't believe it, yet it seems that his assertion is nevertheless acceptable.

In another case, a public school teacher, Stella, asserts to her students that modern humans evolved from other species of primates, on the understanding that that is the view supported by scientific evidence and so the view that should be taught to her students as science. She does not herself believe it however, as it conflicts with her strongly held religious views. She thus asserts what she does not know, but again her assertion seems to be acceptable.

These examples of acceptable assertion without belief, and so without knowledge, involve normatively complex situations. Factors that might be relevant to our sense of what the subjects of these stories are permitted to assert include, not only whatever norm of assertion there may be, but also: (1) Martin and Stella are aware of strong evidence in favor of their assertions, even though they do not think of possessing that evidence as constituting knowledge. (2) They are consciously aware of the dubious or controversial nature of their contrary doxastic tendencies, as in Martin's awareness of his racist feelings, and Stella's awareness that her religious commitment is not shared by her students or her employers. (3) They have duties to the audiences of their assertions, as jurors or teachers or community members (etc.), that may constrain what it is proper for them to say. We have intuitions whether the subjects of these stories should have said what they did, but we do not necessarily

fully understand how we come to have those intuitions. It is reasonable to think that these other factors may influence our intuitive judgments about these cases.

Lackey acknowledges that “role responsibilities” may conflict with the norm of knowledge in some of her cases. But she thinks she has a rebuttal in the example of Martin the racist juror. While acting in the capacity of a juror at the trial, she says, Martin has a duty to set aside his feelings about the accused, and to cooperate in arriving at a verdict compatible with the explicitly enjoined epistemic standard, finding the accused guilty only if it is proved beyond reasonable doubt by evidence introduced in the trial. After the trial is over, she says, Martin no longer has duties as a juror, and so does not have a duty requiring him to set aside the knowledge requirement for assertion.

Is it really so clear however that his duties have ended with his formal participation in the jury process? It seems rather that he might be acting wrongly as a juror and citizen if he says things that undermine the confidence of his community in the verdict of the trial. He may also have a duty not to allow the racist feelings that he now disapproves of to lead him to further damage the reputation of the accused, especially since he believes that detailed evidence tends to exonerate the man.

It seems that Lackey’s cases are not clear counterexamples to the claim that knowledge is the norm of assertion. There is considerable evidence that we have such a norm. So I will continue on the view that we have a knowledge norm for assertion.

Justification, as epistemologists discuss it, is a norm for beliefs, not assertions. How would a knowledge norm for *assertions* lead us to an “appearance of knowledge” norm for *beliefs*?¹⁶ To determine whether we know that *p*, preparatory to asserting that *p*, we must also determine whether we believe that *p*, and whether that belief counts as knowledge. Since it would be difficult to decide for any given belief we acquire whether we are likely to assert it in the future, it would be natural to apply the norm of knowledge to all of our own beliefs that we evaluate. A knowledge norm for assertion will thus in practice lead to a knowledge norm for belief, at least for those beliefs we expect to express (and we have no effective way of anticipating which beliefs those are). The best we can do in practice to comply with the knowledge norm for assertion is to believe that *p* only when there is an appearance to us of our knowing that *p*.

It may still be beneficial to warn someone who can’t tell that her belief is not really knowledge, to prevent her from asserting it. But such correction can’t improve her ability to make appropriate assertions in similar cases in the future. So it may be useful to have a term to mark this distinction, ‘justified belief’. ‘Justified belief’ is a secondary normative term, one that signals whether the subject is adequately trying to comply with norms formulated using the primary normative term, ‘knowledge’.

The knowledge norm of assertion should be easier to teach to new members of society than the related knowledge norms for beliefs and practical reasoning,

¹⁶ In order to make a connection between belief and assertion in a similar context, Douven (2006, p. 453) adopts a suggestion he finds in Williamson (2000, p. 255) and Adler (2002, p. 74), to the effect that belief just is (sufficiently like) internal assertion. The proposal in the text maintains the dissimilarities, in a way that seems to me easier to defend than the assimilation.

because our assertions are available for observation by others, and our beliefs and practical reasoning are usually not.¹⁷ In most cases we will become aware of others' beliefs only when they are expressed in assertions. (The important classes of exceptions to this rule will be when we judge that perceptually normal and apparently alert people probably have correct beliefs about important and easily perceptible features of their immediate surroundings and, in the absence of contrary indications, that they probably have the commonsense beliefs appropriate to people who live in our time and place.) The same goes for practical reasoning and the alleged norm of knowledge for the premises of practical reasoning (Hawthorne 2004; Stanley 2005). Others will usually be aware of the detailed content of our practical reasoning only if we indicate it by making assertions. Norms requiring knowledge for belief and practical reasoning can thus be publicly enforced only as beliefs and practical reasoning become known through assertions. So the norm of assertion must be functionally primary for teaching all of these norms.

Requiring for justification the diachronic appearance of knowledge, rather than the momentary appearance, suggests an objection related to a common argument for an internalism of the present moment about justification (Goldman 1999). It is claimed that justification functions to guide the subject in acquiring and retaining beliefs. But justification could do that only if the subject could be aware of the justificational status of her beliefs, and so could be aware of whatever affects that status. On my view it may be true that someone is justified or unjustified in a belief they hold, even though they don't now remember how they came to hold that belief. So in some quite common sorts of cases they wouldn't be able to tell whether or not they were justified. But how could a belief's being justified, or unjustified, guide their decision in such cases whether to continue to believe it? It seems that any criticism from others will also fail to improve her future performance if she doesn't remember how she came to believe. So shouldn't we endorse an internalism of the present moment and thus reject the diachronic appearance view?

The main reason not to require that justification always be accessible to the subject (in addition to its being incompatible with our intuitions about these cases) is that this requirement makes justification insufficiently accessible to those others who teach and enforce the norms. Public norms must be taught to new members of the community, and reinforced for long time members, through others' assessments of their compliance. It is the others' assessments, and the evidence available for such assessments, that may be expected to determine the contours of whatever normative requirements there may be.

It is the subject's perceptual and testimonial interactions with the world over time that are mainly accessible to observation by others, not what she remembers or believes at present. Suppose Fred is leaving a football game with Sam. Looking at someone walking away in the crowd ahead, he says "Isn't that John? It looks like John. I didn't know he was coming." But they don't approach any closer to the person Fred thinks may be John. A few hours later Sam hears Fred confidently tell

¹⁷ Norms governing belief will thus not immediately apply to assertions and vice versa. For example, a norm of tact applies to assertions, not to beliefs, precisely because the beliefs are not accessible to others and so cannot cause the interpersonal problems that the corresponding assertions would cause.

someone “I saw John at the game today.” It seems that Sam may reasonably correct him, by saying something like “Are you sure? The person you saw was pretty far away and he was walking away from us.” The correction will be reasonable, and may be effective in making Fred more cautious in the future, even if Fred doesn’t now remember the details Sam mentions. Notice that Sam doesn’t claim to know whether or not Fred really saw John: his correction cites facts from which it is reasonable to infer that there was no appearance to Fred as of his knowing that he was seeing John. Sam’s judgment thus concerns whether Fred’s expressed belief could have for him a *diachronic* appearance of knowledge, not whether Fred’s belief was true, or whether it *now* appears to Fred to be knowledge.

It’s not easy to discover what someone else currently remembers of what they once saw, or heard, or thought about. So it makes sense, practically, to assume that others still believe what they perceived or were told, and that they also believe the obvious implications of what they thus came to believe, unless we acquire evidence that they have in fact forgotten or failed to notice what they once saw or heard (etc.). If that sometimes leads us to neglect an opportunity to improve others’ epistemic tendencies, or costs us a bit of energy expressing approval or disapproval where no improvement can be made, that is not likely to be very important in the long run.

Furthermore, in general, it is just not true that we can only learn from corrections of mistakes we remember making. Students may not remember making the mistakes in their papers that I correct, but seeing the corrections still causes them to be less likely to make similar mistakes in the future. There is also a kind of learning in response to success and failure, which continues even where the pattern of cognitively available differences does not rise to conscious awareness. For example gambling behavior over a long sequence of bets will normally be responsive to actual losses and gains, even where the subject is not able to consciously aggregate those outcomes in relation to the different sorts of bets she makes (Damasio 1994, pp. 212–214). We can only be trained to respond to cues that are perceptible, and hence that constitute appearances over time. But we don’t need to process them together, in the “critical review” philosophers like to imagine, to benefit from epistemically relevant corrections.

I conclude that epistemic justification is the diachronic appearance of knowledge. That view explains our reactions to the Gettier cases and also to cases of justified false belief. It is the secondary epistemic norm that would be predicted on the view that we have a knowledge norm for assertion. We do our best to comply with this norm when we assert only beliefs that appear to be knowledge, justified beliefs.

Acknowledgments I am grateful to Arizona State University for a sabbatical leave during which the first draft of this paper was written. Thanks to my students at ASU, to the audience of an American Philosophical Association Colloquium in April 2010, and especially to Stewart Cohen, Bernard Kobes, and Michael White for helpful comments on earlier drafts.

References

- Adler, J. (2002). *Belief's own ethics*. Cambridge, MA: MIT Press.
- Alston, W. P. (1985). Concepts of epistemic justification. *Monist*, 68, 57–89. (Reprinted from *Epistemic justification*, by W. P. Alston, Ed., 1989, Ithaca: Cornell University Press.)

- Alston, W. P. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives*, *II*, 257–299. (Reprinted from *Epistemic justification*, pp. 115–152, by W. P. Alston, Ed., 1989, Ithaca: Cornell University Press.)
- Bar-On, D. (2004). *Speaking my mind*. Oxford: Oxford University Press.
- BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge, MA: Harvard Press.
- BonJour, L. (1998). *In defense of pure reason*. Cambridge, MA: Cambridge Press.
- Chisholm, Roderick. (1991). Firth and the ethics of belief. *Philosophy and Phenomenological Research*, *50*, 119–128.
- Conee, E., & Feldman, R. (2001). Internalism defended. In H. Kornblith (Ed.), *Epistemology: Internalism and externalism* (pp. 231–260). Oxford: Blackwell Publishers. (Reprinted in *Evidentialism: Essays in epistemology*, by E. Conee & R. Feldman, Ed., 2004, Oxford: Oxford University Press.)
- Damasio, A. R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Avon.
- DeRose, K. (1996). Knowledge, assertion and lotteries. *Australasian Journal of Philosophy*, *74*, 568–580.
- DeRose, K. (2002). Assertion, knowledge, and context. *The Philosophical Review*, *111*(2), 167–203.
- Douven, I. (2006). Assertion, knowledge, and rational credibility. *The Philosophical Review*, *115*(4), 449–485.
- Feldman, R. (2001). Voluntary belief and epistemic evaluation. In M. Steup (Ed.), *Knowledge, truth, and duty: essays on epistemic justification, responsibility, and virtue* (pp. 77–92). Oxford: Oxford University Press.
- Feldman, R. (2004). The ethics of belief. In E. Conee & R. Feldman (Eds.), *Evidentialism: Essays in epistemology* (pp. 166–195). Oxford: Oxford University Press.
- Feldman, R. (2008). Modest deontologism in epistemology. *Synthese*, *161*, 339–355.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, *23*, 121–123.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, *73*(20), 771–791.
- Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge*. Dordrecht: Reidel. (Reprinted in *Liaisons: Philosophy meets the cognitive and social sciences*, by A. Goldman, 1992, Cambridge, MA: MIT Press.)
- Goldman, A. (1999). Internalism exposed. *Journal of Philosophy*, *96*(6), 271–293.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, *10*(October), 571–590.
- Huemer, M. (2007). Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research*, *74*(1), 30–55.
- Kornblith, H. (1983). Justified belief and epistemically responsible action. *Philosophical Review*, *92*, 33–48.
- Lackey, J. (2007). Norms of assertion. *Nous*, *41*(4), 594–626.
- Lewis, D. (1996). Elusive knowledge. *The Australasian Journal of Philosophy*, *74*, 549–567. (Reprinted in *Papers in metaphysics and epistemology*, by D. Lewis, 1999, Cambridge University Press, Cambridge.)
- Littlewood J. E. (1986). *Littlewood's miscellany* (B. Bollobas, Ed.). Cambridge: Cambridge University Press.
- Mancuso, K., et al. (2009). Gene therapy for red–green colour blindness in adult primates. *Nature*, *461*, 784–787.
- McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.
- Nelkin, D. K. (2000). The lottery paradox, knowledge and rationality. *Philosophical Review*, *109*(3), 373–409.
- Plantinga, A. (1993a). *Warrant: The current debate*. New York: Oxford University Press.
- Plantinga, A. (1993b). *Warrant and proper function*. New York: Oxford University Press.
- Pryor, J. (2000). The skeptic and the dogmatist. *Nous*, *34*(4), 517–549.
- Reynolds, S. L. (1992). Self-recognition. *The Philosophical Quarterly*, *42*(167), 182–190.
- Reynolds, S. L. (1998). Evaluational illusions and skeptical arguments. *Philosophy and Phenomenological Research*, *58*(3), 529–558.
- Reynolds, S. L. (2002). Testimony, knowledge and epistemic goals. *Philosophical Studies*, *110*, 139–161.
- Smithies, D. (January 2011). The normative role of knowledge. *Nous*. doi:[10.1111/j.1468-0068.2010.00787](https://doi.org/10.1111/j.1468-0068.2010.00787).
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Sutton, J. (2007). *Without justification*. Cambridge, MA: MIT Press.
- Unger, P. (1975). *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.