

Transferring Knowledge

PETER J. GRAHAM
Saint Louis University

Most of what we know we know because we accept the word of others. Why is communication or interlocution—testimony in the broad sense—a source of knowledge? Why do I sometimes learn things about the world by talking to you, by accepting what you say? The standard answer is that communication is a source of knowledge because the knowledge that a speaker possesses is somehow *transferred* to the hearer. The standard view holds, among others, the following two claims. The first is a necessity claim; the second is a sufficiency claim.

- (KN) H comes to know that P by accepting S's statement that P only if S *knows* that P.
- (KS) If S *knows* that P and S sincerely states that P and H *justifiably* accepts S's statement that P then H comes to know that P.¹

I argue against this view and in favor of two replacement principles. The information-theoretic alternative—an account pioneered for perceptual knowledge by Fred Dretske—endorses the following two claims.² The first is a necessity claim; the second is a sufficiency claim.

- (IN) H comes to know that P by accepting that P only if H's basis for accepting—H's internal, cognitive state of understanding S as having asserted that P—*carries the information* that P.
- (IS) If H's basis for accepting that P—H's internal, cognitive state of understanding S as having asserted that P—*carries the information* that P and H *justifiably* accepts that P then H comes to know that P.

All four principles use the locution 'accept'. Acceptance involves understanding a speaker as having asserted or presented-as-true (as having stated or reported) that P and then believing that P on the basis of so understanding.³

Acceptance is a species of belief; it is not meant to mark out a cognitive state that is somehow distinct from belief.⁴

In this paper my primary aim is undermining (KS) and supporting (IN), though I will say a word or two in support of (IS). §I concerns what it is for something to carry the information that P and why information carrying should be relevant to knowledge. §II-III is the heart of my case against (KS) and in favor of (IN). §IV treats a modification of (KS) due to C.A.J. Coady in his important and highly praised book, *Testimony: A Philosophical Study*. §§V-VI discuss two objections and §VII raises a general issue. I take up the argument against (KN), and further defense of (IS), elsewhere.⁵

I

I endorse the following necessary condition upon knowledge: to know that P one must base one's belief that P on adequate grounds. Adequate grounds that P are those that *establish the fact* that P or *guarantee the truth* that P (Dretske 1971, 1981, Armstrong 1973, Fogelin 1994, McDowell 1994). Grounds guarantee the truth that P just in case the grounds would not be as they are unless it were the case that P. What makes the subjunctive true are the operative laws *in situ*. What the grounds guarantee is thus understood in terms of nomological dependencies. One way to understand all of this is in terms of *information carrying*.

Information carrying is due to a law-like correlation or counterfactual dependence between a signal—an event, condition, or state of affairs, including such things as utterances and mental states—and another event, condition, or state of affairs. The rings of a tree, for example, carry the information or *indicate* the age of the tree. Footprints of a certain sort in the snow carry information about wildlife in the vicinity.⁶ All of this, however, is *relative* to circumstances, to the local conditions that obtain. When a signal carries the information that P it is a guarantee *in the circumstances* that P. Ordinarily certain tracks or marks in the snow would not be there *unless* local wildlife were there. But if the local fraternity placed tracks in the snow for fun, tracks of the same type as quail tracks, then the possibility that *these* marks right *here* were caused by a college frosh and not by a quail is now a *relevant alternative possibility* to the proposition that they were caused by wildlife. At best, in these circumstances, the marks in the snow carry the *disjunctive* information that wildlife *or* pranksters are present in the surrounding forest.

What has this got to do with knowledge? Perceptual knowledge that P requires that the subject base her belief on a perceptual or cognitive state that carries the information that P. Why should this be so? The information-carrying requirement incorporates the popular “no relevant alternatives” requirement, and thereby solves Gettier and post-Gettier cases.⁷ Here the idea is that if I cannot perceptually distinguish P from Q—P and Q are “perceptual equivalents” or tokens of the same type (Goldman 1976)—and the possibility that Q is *not* relevant, then the fact that I cannot distinguish between P and Q does not prevent me from percep-

tually knowing that P. The possibility that I am deceived by an evil demon does not rob me of my knowledge that I am sitting here by the fire, for such a possibility is not a *relevant* possibility. But if the possibility that Q is relevant, then perhaps I do not know. The shopworn case is where someone is driving through the countryside and sees what she takes to be a barn that is in fact a barn, but all around in the surrounding countryside are barn façades. The driver cannot tell whether what she sees is really a barn or just a façade. The possibility that what she sees is not a barn but a façade is relevant and since she cannot rule out the possibility from her car, it seems she does not *know* it is a barn (Goldman 1976). *In these circumstances* her perceptual state does not carry the information that a barn is on the hill, for a token perceptual state of the same type could also be caused by something that is not a barn.⁸ That is why she does not know. This approach has the dual advantage of capturing the sense in which it is true that if you know that P then you cannot be mistaken, the sense that is stronger than the mere triviality that knowledge is belief that is true, in that it requires true belief in actual and nearby possible circumstances, and of capturing the sense in which knowledge is not infallible belief, belief that cannot be mistaken regardless of circumstance.

To summarize, knowledge requires adequate grounds. Adequate grounds establish the fact or “track the truth.” Truth-tracking is a subjunctive property of the grounds. The subjunctive holds because of the operative laws *in situ*, because of the lawful dependencies between the grounds and the facts. So to know that P one must base one’s belief on grounds that carry the information that P.⁹

Stating this as a necessary condition, however, does not exclude the importance of *justifiably* believing that P as a component of knowledge, *pace* Dretske. However, I will not, at least in this paper, attempt an *account* of justifiably believing.¹⁰ This is not to say that the issue will not arise. It will. But when it does I will only say a few words about what justifiably believing something on the basis of what another person says does *not* require.

If knowledge that P is belief that P justifiably based on grounds that carry the information that P, it follows that in the case of belief based upon understanding a speaker’s utterance as the report or presentation-as-true that P, for the hearer to know that P her internal, cognitive state of understanding the speaker as so presenting must carry the information that P. What follows is (IN). And if the hearer justifiably believes that P in such a case, she will know that P. What follows is (IS).¹¹ However, though I endorse the general claims that I have made and so believe that (IN) and (IS) follow as special cases, I do not rest my argument in favor of these principles and against the standard view on the general claims. Rather, my argument will turn on the examples presented in §§II-III and the defense carried out in the sections that follow.

I want to understand why we sometimes acquire knowledge by accepting what other people say. The standard answer, as I see it, is largely metaphorical in inspiration; it relies on the metaphor of *transferring goods* from one person to another. If person A does not have good G, then person B will not come to have

good G either by accepting A's offer. And if A does have G, then B will come to have G too if B accepts A's offer. But when it comes to knowledge, I do not see how it is that *knowledge* gets *transferred*. I do not see how the standard view *explains* why communication is a source of knowledge. Something else must be going on. I return to this larger, explanatory question when concluding.

My thinking on these matters is driven largely by reflection on a certain range of cases, the range of cases that the information-carrying approach to knowledge accords well with, primarily perceptual ones. The cases I will discuss involve a subject who knows that P because he or she has seen, smelled, tasted, or heard that P, or has perceived some relevant fact and then inferred, on the basis of a general belief, that P, or at least would so infer on some occasion. I will not discuss cases where the subject's report is not about a particular, perceptible matter of fact. So I will not discuss reports about general matters of fact, or about moral matters, or about mathematical matters or matters of necessity. It may be that the truth of my view is limited to perceptual or observational reports, or perhaps also reports of introspection, and that there are no possible counter-examples to the standard view when the reports involve mathematics, general facts, and so on. Here I leave these questions open.

II

I will offer two counter-examples to (KS) that support (IN).¹² The first case is a variant of an example due to Harman (1973) and first introduced into the testimony literature by Adler (1996). Call this the *Newspaper Case*.¹³

The military of a small country hopes to stage a successful coup, and pays off or threatens the reporters of the country's newspapers to report that the President has been assassinated regardless of what in fact happens. All but one of the reporters gives in. Andy will report what really happens, and not just what the military wants him to report. As it turns out, the assassination attempt is successful and Andy is the only eyewitness. The other reporters do not know or even inquire into what really happened. Andy writes in his by-lined column that the President was assassinated. When Jenny reads Andy's article, does she know that the President was assassinated?

If we suppose that Jenny has no special relationship to Andy or his paper, and would rely on any of the other newspapers if Andy's were not available, or if she just felt like reading any old paper, or were even disposed to occasionally glance and rely upon the headlines of the other papers, then I think it is clear that Jenny does not come to know that the President was assassinated even though Andy, the person whose report she is accepting, does know that the President was assassinated. The case thus undermines (KS).¹⁴

It is useful to think of this case on analogy with the barns case. In that case things would "look" the same to the subject even if she were to look at a barn

façade. The barns she sees and the façades she might have seen are perceptually equivalent; they cause experiences of the same type. Though they would not look exactly alike, indeed the real barn and the fakes may look very different, e.g. the real barn is old and small and the fakes are new and rather large, either look disposes her to form the belief that there is a barn on yonder hill. She can distinguish the individual barns and barn façades from each other *as individual objects*, but she cannot distinguish the barns *as barns* from the façades *as façades*. Relying on perceptual appearance alone, she cannot tell whether something is real or a fake. Each look, for her, plays the same cognitive role, disposing her to believe there is a barn. And since, in the circumstances, there are such fakes, seeing a barn does not carry the information that there is a barn. The same kind of experience would be caused by something that is not a barn, so the experience of seeing a barn does not guarantee that there is a barn.

Now suppose we think of Andy's report as a "stand in" or "surrogate" for the fact that the President was assassinated, for his report *guarantees* that the President was assassinated, and so is "just as good as the real thing." And suppose we think of the other reports by the other reporters as "imposters" or "fakes." We can then see the analogy between Jenny and the subject driving through the countryside: Andy's report is *like* the barn and the other reports are *like* façades. Jenny, though she reads a report that is just as good as having the fact before her, does not come to know because she would have formed the same belief on the basis of the same type of cognitive state, reading a "fake" report. Andy's report and the other reports are, for her, cognitively equivalent. Though perhaps the reports look very different, e.g. they are printed in different fonts with different bylines and report different details, and so on, they both dispose her to believe that the President was assassinated. Each potential state of understanding the various reports plays the same cognitive role, which is evidence that they are, for her, of the same type. Since Jenny does not know when she relies on Andy, a knowledgeable reporter, the case clearly undermines (KS).

Does the example support (IN)? Is it a case where the information that P is lacking? One might say that it is not such a case. One might say that the information that P is present in Andy's report, for Andy saw the assassination and would only report what he saw. Indeed, since the information is present and Jenny does not come to know, it might be said that the example not only does *not* support (IN), but that it also *undermines* (IS). Both suggestions, however, would be based on a mistake. What matters according to (IN) and (IS) is the hearer's internal, cognitive state of understanding the speaker as having made a certain report, and whether that state carries the information that P. If the speaker's report carries the information that P it does not follow that the hearer's internal state of understanding the speaker carries the information that P. Consider the following example. If lightbulb A is on, then given the wiring, B is then turned on, and B in turn turns on C. B's being on carries the information that A is on, and C's being on carries the information that A and B are on. Now suppose C's being on will result in D's being on, but E's being on will also bring it about the D is on. When D is

on, that means (carries the information) that either A (and B and C) is on *or* E is on. In Jenny's case, Andy's report (C) carries the information that the President was assassinated (A) (and that Andy believes that, (B)). What about Jenny's state of understanding Andy's report (D)? Does it carry the information that the President was assassinated (A)? Since the other reports (E) would produce a cognitively equivalent state of understanding, the answer is no. At best her state of understanding Andy's report carries the disjunctive information that there are fake reports *or* the President was assassinated. Andy's report carries the information that P, Jenny's grounds for believing that P do not. The case supports (IN) and does not threaten (IS).¹⁵

III

In the case just discussed, the reporter was reliable in the relevant sense, but the presence of unreliable reporters prevented those who rely upon him from acquiring knowledge. In the following example the speaker is not reliable in the relevant sense because the speaker sometimes makes a false higher level classification of objects and events in his or her experience, because of a *false connecting belief*. The example is a variant of a case due to Dretske (1992). Call this the *Winetaster Case*.

George is a connoisseur of fine wines. He unerringly identifies a Médoc wine as a Médoc when he tastes one. He knows that Médocs are Bordeaux, for he knows that Médoc is a region of Bordeaux. George also unerringly identifies Chianti wines as Chianti. Hence he unerringly distinguishes Médocs from Chianti on the basis of their respective tastes. However, strangely enough (see note 17), George mistakenly identifies Chianti as Bordeaux, for he mistakenly thinks that Tuscany is a wine growing region in southern Bordeaux. At a dinner party where both Médoc and Chianti wines are being served, George is served a Médoc and correctly identifies it as such, and thus knows that it was Bordeaux. The next day his friend Michael asks him what kind of wine was served and George tells him it was Bordeaux.¹⁶

The various conditions required by (KS) are satisfied, but he does not know what he was told. Hence (KS) is false.

George's testimony does not put Michael in a position to know that the wine was a Bordeaux because George's testimony does not carry the information that the wine was a Bordeaux. At best his testimony carries the information that the wine was a Bordeaux *or* a Chianti. In these circumstances George would say that P whether or not P. When Michael comes to believe that the wine was a Bordeaux, he may form a true belief, even a justified true belief, but not a belief that counts as knowledge. The case thus supports (IN). Structurally what is going on in this case is that the relevant alternatives that the speaker can rule out, or has ruled out,

are greater than those the hearer can rule out just by relying on the speaker. The speaker has failed to “immunize” the hearer from all of the relevant alternatives.

Many do not find this example convincing as a counter-example to (KS) at first pass. The most common reason offered is that George does not know because of his false geographical belief. Is this reason persuasive? Well, it depends on how we flesh out the case. Suppose George thinks every red wine from Europe is from a region of Bordeaux. Then George does not know that wines from Bordeaux are in fact from Bordeaux. But suppose that George is only wrong about Tuscany. Shouldn't we say then that George *does* know that Médoc is a region of Bordeaux? Wouldn't George be like a typical New Yorker who thinks New Mexico is in Mexico but still knows that Wisconsin is in the United States? A typical New Yorker isn't guessing when he says Wisconsin is in the United States. Surely we can say that George knows, despite his false belief about Tuscany, that Médoc is a region of Bordeaux.¹⁷ Furthermore, the fact that Tuscany is not a region of Bordeaux is not a *defeater* for George's belief that the wine he drank was a Bordeaux, for if he were to come to know that Tuscany is not in Bordeaux he would still know that Médoc is in Bordeaux and would still believe that the wine, which he correctly identified as a Médoc, was a Bordeaux.

There are two other reasons that may persuade one to reject the example. The first is due to Coady's discussion of the example.¹⁸ Coady's thinks George's evidence is really no better than Michael's evidence. Coady first states the analysis I favor:

It may be urged that there is a *crucial difference* in the status of the evidence that George and Michael have for their common belief that the wine was a Bordeaux. George's evidence is that it tasted like a Médoc and he knows that all Médocs are Bordeaux, so it is plausible to think that he would not have this evidence for his belief unless his belief were true. Michael's evidence is that George says that the wine was a Bordeaux and this we might suspect is evidence he would have even if his belief that it was a Bordeaux were false, i.e. in the circumstances that it was a Chianti. (1992: 227–8. Emphasis added.)

Why does Coady think this is wrong? He says that:

It is not clear how we can treat as true the claim, ‘it is possible that Michael should have had this very evidence but his belief should have been false’, without making it true also that ‘it is possible that George should have had this very evidence but his belief have been false’. We could treat Michael's evidence as merely some utterance of George's to the effect that some wine he had at a dinner party was a Bordeaux, but then why not treat George's evidence equally abstractly as some taste sensation which he identifies as belonging to a type of wine which he believes to be a Bordeaux? To do so is to allow that in both cases the evidence might remain the same and the belief be false. (1992: 228)

Coady's objection, in essence, is that we should treat George's evidence such that there is *no difference* between his coming to believe that a *Chianti* is a Bordeaux

and his coming to believe that a *Médoc* is a Bordeaux. There is no clear reason, Coady appears to suggest, to hold that George's evidence is better than Michael's once we re-describe the case or treat George's evidence "equally" abstractly as we treat Michael's evidence. George's evidence is thus "tastes like a Bordeaux" when it is Bordeaux *and* when it is Chianti. George's evidence is thus no better than Michael's is.

This will not do. What does it amount to other than the suggestion that if we abstract away from the relevant features of the example in question that make it work as a case against principle (KS) that it will no longer work? The evidence that George possesses is not "tastes like a Bordeaux" because Médoc wines and Chianti wines produce *different* taste-sensations for George. He accurately and reliably distinguishes Médoc from Chianti on the basis of their different tastes. George's mistake is to group Médocs and Chianti together as Bordeaux on the basis of a false *geographical* belief; he does not do so on the basis of some "higher level" taste sensation. Indeed, it is consistent with the example to suppose that George thinks that Chianti taste very different from Médocs, and from every other Bordeaux.¹⁹

What the example shows is that when the speaker can rule out more relevant alternatives than the hearer, a speaker can sometimes know that P and sincerely state that P but not enable his hearer to learn something on the basis of his testimony. This is because the speaker's *belief*, though it is caused or causally sustained by a signal that carries the information that P, does not *itself* carry the information that P; and it does not carry the information that P because, in this case, the speaker has a false general belief that sometimes causes him to believe that P when not P.²⁰ Because his belief does not carry the information that P, when he states that P his statement will not carry the requisite information either. Hence he will know something even though he cannot bring his audience to know it by communicating the belief to them.

The only other reason I can imagine for finding the example unconvincing is that it contradicts (KS). It is true that (KS) is intuitive. But I do not see that it is unassailable, nor do I see that its plausibility is greater than the plausibility that George knows and Michael does not. Further, the plausibility of the principle, I claim, is due to the fact that it is generally true that if a speaker's statement carries the information that P, that is because the speaker's belief that P is caused or sustained by the information that P. Hence it will generally, though not always, be the case that a speaker enables a hearer to come to know that P when the speaker, in fact, knows that P. The plausibility of (KS), then, is due to the truth of (IN). I return to this point below.

IV

There are two ways to treat the counter-examples without endorsing (IN). The first is to concede the force of the examples without giving up (KS) as expressing a sufficient condition on the speaker's side of the equation, instead holding that

the example shows that (KS) should be modified to add further conditions to the list of what the *hearer* must do for knowledge to be transferred from the speaker to the hearer. The second is to concede the force of the examples but hold that (KS) should be modified; further conditions should be added to the list of what the *speaker* must do to transfer knowledge. Coady, for one, takes this latter route. First I discuss Coady's view as an exemplar for views that add conditions upon the speaker and then I briefly discuss views that add conditions upon the hearer.

I will first take up the Winetaster Case and then discuss the Newspaper Case. When concluding his discussion of George and Michael, Coady suggests that (KS) ought to be reformulated to add a "competency" condition. Once he adds such a condition he *can* concede that George knows about the wine even though George cannot transmit that knowledge, for he is not competent about Bordeaux wines (1992: 229–30). So knowledge is not sufficient, but knowledge plus "competence" is. Hence Coady would endorse (KS*).

(KS*) If S knows that P, and S has the relevant competence, authority, or credentials to state truly that P, and H justifiably accepts S's sincere statement that P, then H comes to know that P.²¹

The problem (KS*) faces is stating exactly what is to count as the "relevant competence, authority, or credentials to state truly that P." Naturally we should beware of explicating "competence" or "authority" or "credentials" as "whatever is needed beyond what (KS) requires to transmit knowledge." That would not be very informative. But once we explicate what competence is, the proposal faces a dilemma: either George is competent, in which case the example is still a counter-example to the new principle, or George is not competent because his testimony does not carry the information that P. In short, either (KS*) is of no help or (KS*) is just a version of the information-theoretic account.

To take up the first horn, consider what would be the most ordinary reading of "the relevant competence, etc." That reading suggests that George *does* possess the relevant competence, authority, or credentials to state truly that P. That is, suppose that George knows almost everything there is to know about European wines. His one mistake is the belief that Chianti is produced in a region of Bordeaux. If he were called in to a courtroom as an expert witness about European wines he would most surely pass muster (see note 17). A fact about competence in any skill or subject matter is that it does not require perfection or infallibility.²² So when he truly states that the wine was a Bordeaux, doesn't he pass Principle (KS*) as well as Principle (KS)? But knowledge is not transmitted, so (KS*) is false as well.

Now suppose we understand "the relevant competence, etc." in a more technical way, in a way that will avoid the counter-example. What Coady needs is an account of competence that is relative to particular propositions *and* is such that if a speaker is competent relative to P then it is guaranteed that P. That is, what Coady really needs to add is *not* that the speaker is competent (in the ordinary

sense of 'competent') with respect to P, but that his statement that P carries the information that P. At best (KS*) is really just (KS) plus (IN), and (IN) is doing all the work.²³ (KS*) offers no real alternative to (IN) and should be abandoned along with (KS).

It is even easier to see that (KS*) is no help with the Newspaper Case. Andy is clearly competent, in both the ordinary sense and in the technical sense, for he saw the assassination and he is an honest reporter. However, it is also clear that competence in either sense is not enough, for Jenny does not learn that the President was assassinated when reading Andy's article. We can now see that (KS*) is not even (KS) plus (IN), for (KS*) has to do with the speaker and his or her statement, and not with the hearer's grounds, with the hearer's cognitive state of understanding the speaker.

What conditions on the hearer could one add to explain the case? What is required is a general condition that a hearer must meet, a condition that the hearer fails to meet in the cases described. First, one could require that the hearer *know* that the *speaker know*. Michael surely does not know that George knows, and Jenny does not know that Andy knows. But this is clearly too demanding. The whole point of relying on others is that one can acquire knowledge from another where one knows less than one's interlocutor, or at least less about the particular case. Now one might object by pointing out that the hearer need not know *how* the speaker knows in order to know that the speaker knows: the hearer can possibly know that the speaker knows in some independent way, perhaps in virtue of knowing that the speaker is reliable because the hearer has checked up on the speaker's reliability for herself or because another person has told her so. There are two problems with this. First, knowing that the speaker is reliable is arguably too demanding, and knowing from another's testimony begs the questions at issue. Second, to know that the speaker knows in an independent way would require knowing that P in an independent way, for P is what one will know if one knows that someone knows that P. But once the hearer knows that P in an independent way, why should the hearer bother to rely on the speaker?²⁴

Second, one could require that the hearer *know* that the speaker is *trustworthy, justified, or reliable*. But that would not work as the considerations just raised show: it also would be too demanding. Further, isn't George trustworthy and reliable, at least generally? Though he is not reliable about the geography of southern Bordeaux, isn't he reliable about wines? If we say he is not trustworthy or reliable because he sometimes gets it wrong about Bordeaux, then to save (KS) we must add that the hearer knows that the speaker is reliable in the sense of ruling out the relevant alternatives in the particular case. But again, that is asking too much of an ordinary hearer. (See the fourth proposal, below.) Lastly, such a move is contrary to the spirit of the idea that one can learn from others by trusting them, for requiring that one know that one's source is reliable is not really to trust one's source but rather to trust oneself.

Third, one could require that the hearer be *justified in believing* that the speaker *knows*. But again this would not work for it seems that Michael and Jenny would

both be justified in holding such a belief, and so would satisfy such a condition. Also, such a condition is arguably too demanding, at least for children and for ordinary cases of interlocution.

Lastly, one could require that the *hearer* rule out the relevant alternatives that the speaker has not ruled out. That would do the trick, but that would amount to endorsement of (IN), for then the hearer's state of understanding the speaker, given what else the hearer knows, would carry the information that P.²⁵

I see no other condition that would work. For a further condition on the hearer to work it must either be too demanding, or really just be the endorsement of (IN). So adding conditions on the speaker or the hearer will not save (KS) in letter or in spirit.²⁶

V

I now offer a possible objection that involves both (KS) and (IN). The boy who cries "Wolf!" is not to be trusted. The members of his community have been fooled one too many times. But then one day he *sees* a wolf and thereby comes to *know* that there is a wolf. He runs up the hill and encounters a stranger who is unaware of the boy's previous false steps and cries "Wolf!" The stranger, alarmed, accepts the boy's report. Does he come to know that there is a wolf threatening the village herd?

A hypothetical objector might hold that this case poses a difficulty for my view.²⁷ One might say that the stranger comes to know because the boy knows, even though the boy is not a reliable reporter about the presence of wolves because of his past false reports. The boy's report does not indicate the presence of a wolf because he would say P even if not P. Nevertheless he imparts knowledge because he "speaks from his knowledge." (KS) is true, the hypothetical objector claims, and (IN) is false.

It is unclear exactly what to say about this example without spelling it out in more detail. There are at least two ways to do this. In the first it is plausible to hold that the stranger *does* come to know and in the second it is clear that the stranger does *not*. Suppose that the boy, who has cried "Wolf!" a number of times in order to get attention from the villagers, learns his lesson at the very moment when he actually sees a wolf. He is so scared that his psychology takes a turn for the better. He now says that there is a wolf *because* he sees one. He realizes the gravity of the present situation, and he even realizes how terrible his previous deeds were. Before he was, in a way, a broken alarm. He would go off when no danger was present. But now that he detects danger, his report of the threat to the village is a reliable one. The fright is sufficient to fix the alarm. *In these circumstances* his statement indicates the presence of a wolf. And so it is not so unreasonable to say that the stranger, when accepting the boy's report, comes to know that a wolf is threatening the flock.²⁸

But we can also spell out the case in another way. Suppose the boy still cares little for honesty, but it is not the attention he is after, but rather the sheer delight

in making reports about important matters in a completely random way, that leads him to tell the stranger that a wolf is present. That is, seeing the wolf prompts him to make a report about the presence or absence of a wolf, but the reason why he says a wolf is present and not absent is because he flipped a coin. He says there is a wolf not because he believes on the basis of excellent evidence that a wolf is present, but because the coin turned up heads. He may be a reliable *believer*, but he is not a reliable *reporter*. Here the stranger does not come to know about the wolf. He is in no better shape than he would be in if the boy had asked him if he thought a wolf were present and he made up his mind by flipping the coin himself.

What the example shows is that it is important not to think of reliability in *merely statistical* terms, for if we think in those terms, we will say that the boy cannot be a reliable reporter, even when he sees the wolf and is shocked into a change of disposition, for most of the boy's reports about wolves are misleading. What matters is the truth of the subjunctive on particular occasions. What matters is whether, on this occasion, the boy would not say that there is a wolf unless there is a wolf. So even though the boy has made more false reports about wolves than true ones, one can still learn about whether there is a wolf threatening the flock from the boy once he has seen one before his very eyes. What matters, in short, is the disposition underlying the boy's report and whether it will underwrite the truth of the subjunctive. It is this mistake, thinking of the relevant reliability in statistical terms, that led the hypothetical objector into thinking that what matters is whether the speaker knows that P, not whether the speaker's report carries the information that P.

Our hypothetical objector is wrong to think that the stranger can come to know even when the reporter is relevantly unreliable, even when the speaker's report does not *indicate* that P. When it is clear that the stranger comes to know it is because the boy is *now* a reliable reporter. He *now*, or at least *in these circumstances*, says that P only if P. The boy *was* an unreliable reporter, but he no longer is. He is like a thermometer that once was broken but now is fixed. Although the case does not impugn (KS), for when the hearer does not know the report is not relevantly sincere, it does not impugn (IN) either.

VI

My thinking on these matters has been guided by the following idea: the epistemology of testimony is importantly analogous to the epistemology of perception. In this respect I follow Thomas Reid.²⁹ But it may be thought that there is an important *disanalogy* between testimony and perception. It is reasonable to believe that our perceptual experiences and beliefs carry information because there are laws or law-like regularities that (in normal circumstances) correlate our experiences and beliefs with the facts. But is it reasonable to believe that something like this obtains in the case of human testimony? Is testimony really like perception? After all, in the case of perception and perceptual belief, nature is doing *most*, if not *all*, of the work. But when it comes to testimony, *agency* is involved.

The speaker *chooses* to say what he or she does; it is up to her to say what she believes. Testimony, unlike perception, involves another mind. More importantly, it involves *free will*. The speaker *can* always choose to *lie* or *mislead*. Hence there is an important disanalogy between perception and testimony.³⁰ Perhaps it simply does not make sense to say that a speaker's utterance or report is something that, like the rings of a tree or marks in the snow, carries information about some other event or state of affairs. The reliability of testimony, it might be said, is not something that is underwritten by laws or law-like regularities *in situ*.

I think we sometimes know things on the basis of perception because of the circumstances we are in, the laws of nature, and our cognitive makeup. The circumstances, the laws, and our psychology underwrite, when true, the subjunctive that things would not look like this unless such and such were the case. Though it is true in one sense of 'can' that things can always go wrong, that things might always look differently than they are (a demon might be fooling us, the laws of nature might have suddenly changed, and so on), it is *false* in another sense that things can go wrong *given* the circumstances, the laws, and the cognitive machinery in question. A perceptual experience carries the information that a tree is before one *given* that lighting conditions are normal, that one's eyes are functioning properly, and so forth. *Given* all of that, it is *false* that there might not be a tree but it seems as if there is a tree. All counterfactuals, all subjunctives, presuppose that certain actual conditions, conditions that are independent of the imagined change, *remain actual* in the envisaged counterfactual situation. When we imagine changes to test the truth of the subjunctive, we hold fast the laws of nature, my cognitive makeup, and so on.

When it comes to testimony I think we sometimes come to know things on the basis of what other people say because of the circumstances, the laws of nature, and the cognitive makeup of our interlocutors and ourselves. On the speaker's side, we learn things (in part and usually) because speakers would not say that P, for some P at some time in some circumstance, unless they believed that P. For instance, you know that I had waffles for breakfast today in part because I would not say that I had waffles today unless I *believed* I did. I am that kind of person. When we test the truth of this subjunctive in a counterfactual situation we hold fast traits of character, general motivations, and other facts that are independent of what I had for breakfast. We are saying that *in those conditions*, the conditions that actually obtained on the occasion in question, one event regularly follows another. Compare not just what I say, but also what I do or what happens to me. Walking down the street it is highly unlikely that I will trip and fall. I am that kind of person. I am not a clumsy ox. But if Chester were to trip me, then I would fall. When we say "I would not have fallen if Chester had not have tripped me" we assume that I am the kind of person who does not regularly stumble over his own feet. If this condition is not held constant, then the subjunctive will turn out false. I would or might have fallen even if Chester didn't trip me. In evaluating subjunctives or counterfactuals, we hold certain conditions constant, whether the

subjunctives have to do with marks in the snow, the rings of a tree, the production of perceptual experiences, linguistic behavior, or falling down face first.

What makes it true that speakers would not say that P, at least for some P at some time in some circumstance, unless they believed that P? Presumably it has to do with the *character* of the speaker in question, with the speaker's cognitive makeup, with his or her *psychology*. Reid held that we are naturally disposed to say what we believe. He called this natural disposition the *principle of veracity*, an innate or original principle of the human mind.³¹ Reid, I think, was correct to posit such a disposition. Are there other reasons that could explain why speakers would (for some P, for some time, and for some circumstance) not say that P unless they believed that P? Presumably there are. It may be due to upbringing. It may be due to a respect for the duty to tell the truth. It may be due to hypnosis or mental handicap. It may be due to the belief that if one is honest one will enjoy great rewards in the next life. People are honest because they are naturally inclined to honesty and have no reason to mislead, or because they were raised that way, or because they believe it is the right thing to do, or because they are compelled by another or by mental handicap, or out of self interest. Explanations of behavior that appeal to beliefs, desires, innate principles of reasoning, and so on, are *intentional* explanations. They are explanations that appeal to intentional or psychological *laws*. Compare walking to the fridge for a beer and saying that there is a beer in the fridge. Even though ordinary behavior, such as walking into the kitchen to get a beer out of the fridge, involves agency, there are still psychological or *ceteris paribus* laws underwriting the production of the behavior. He wants a beer, thinks there is a beer in the fridge, and walks to the fridge and opens the door because of what he believes and desires, *ceteris paribus*. Such *ceteris paribus* laws are laws that involve facts about *character* or *individual psychology*, laws that will underwrite, when true, the subjunctive that a speaker would not say that P unless he or she believed that P.

Psychological laws, though perhaps not reducible to laws of nature, are nonetheless *laws*. This was a point Hume was well aware of.³² It is a point we exploit when reasoning about what other people will do so that we can coordinate our actions, or about what they believe and desire so that we may bend their behavior to accord with our own beliefs and desires.³³ When we reason about actions and natural effects, we treat them on a par. Hume wrote, of a particular example, that:

Here is a connected chain of natural causes and voluntary actions; but the mind feels no difference between them in passing from one link to another: Nor is less certain of the future event than if it were connected with the objects present to the memory or senses, by a train of causes, cemented together by what we are pleased to call a *physical* necessity. The same experienced union has the same effect on the mind, whether the united objects be motives, volition, and actions; or figure and motion. (*Enquiry Concerning Human Understanding*, VII, I, 70.)

Even if there is an important difference in the nature of things between physical necessity and action, it does not follow that there is an important difference

in the regularity with which physical laws and psychological laws produce the motions of the wind and the voluntary behavior of agents. Even though testimony, unlike perception, involves another mind it does not follow that there is an important disanalogy between perception and testimony. It does not follow that it does not make sense to say that a speaker's utterance or report is something that, like the rings of a tree or marks in the snow, carries information about some other event or state of affairs. Again there are two senses of 'can' at work here. In one sense the speaker can always choose to lie or mislead. But in another sense, given the speaker's character, given what else the speaker believes and desires, and so on, it is *false* that the speaker can always choose to lie or mislead. "I cannot tell a lie" a famous American once said, or so Americans are fond of saying. I know I might choose to say that I had something else for breakfast, but I would not say that I had waffles unless I did. Really, I *am* that kind of person. Though the explanations, at one level of description, for why perception and testimony carry information that cognizers pick up an utilize may differ, it does not follow that one provides information and the other does not. If there is an important disanalogy between testimony and perception, a disanalogy that undermines the information-theoretic approach to testimony, the existence of free will and the possibility of deception is not its source.³⁴

VII

I said I wanted to understand why testimony is a source of knowledge, why it is we sometimes come to learn something by accepting the word of another. Although I have only argued here in favor of (IN), I think the reason we sometimes learn that P from others is because the reports speaker's make carry the information that P, and that information is information that, in the ordinary case, we will pick up and make use of when coming to believe that P. Let me fill this out in a little more detail, and at the same time try to account for the plausibility of (KS) and (KN).

I claimed at the outset that knowledge that P is belief that P justifiably based on grounds that carry the information that P. In the ordinary case when a speaker knows that P, his or her belief that P will not only be based on grounds that carry the information that P, but the belief itself, because it is based on such grounds, will also carry the information that P. Although this does not obtain in George's case, it does obtain in Andy's case, and it will obtain in most cases of ordinary perceptual knowledge. Your eyes, for instance, will not only produce a perceptual experience that "tracks the truth," the belief you form will also track the truth.

Furthermore, in an ordinary case where a speaker says what he or she believes, his or her belief will be a part of the cause of what he or she says. As a result, the assertive utterance, the report, will also carry the information that P. In the ordinary case, then, when a speaker knows that P and sincerely states that P, the statement produced will carry the information that P.

Lastly, in an ordinary case, a hearer's cognitive state of understanding a speaker as having asserted that P will be the result of a speaker having asserted that P. If the speaker's assertion or statement carries the information that P, then in the ordinary case, the case where there are no relevant alternatives, the hearer's state of understanding the speaker will also carry the information that P.

In short, in the ordinary case the information that P flows from the basis of the speaker's belief to the speaker's belief to the speaker's statement to the hearer's basis for his or her belief, and then to the hearer's belief, which in turn may flow to the hearer's statement when the hearer passes the message on to another interlocutor.

What all of this shows, I think, is that in the ordinary case a hearer learns from a speaker not because the speaker knows, but because in knowing that P the speaker believes that P on the basis of an internal state that carries the information that P, information that he or she can pass on to a hearer through communication.³⁵

These facts explain the plausibility of (KS), for ordinarily when the speaker knows the information that P will be conveyed to the hearer.³⁶ This also explains the plausibility of (KN), for when the speaker does not know, that is because, at least ordinarily, the basis of his or her belief does not carry the information that P, and so the information is not conveyed to the hearer. But when the speaker's *belief* does not carry the information, or when the speaker's *statement* does not carry the information, or when the hearer's *cognitive state* of understanding the speaker as having stated that P does not carry the information, then even though the speaker knows that P, the hearer will not come to know that P by trusting the speaker. Even though they are plausible, they are false. A hearer needs the information that P to know that P, something a knowledgeable speaker may fail to provide.³⁷

Notes

¹For endorsements of (KN) see Audi 1997, Welbourne 1994, Dummett 1994, A. Ross 1986, Burge 1993, J. Ross 1975, Lehrer 1987, 1990, 1994, Hardwig 1985, 1991, and Plantinga 1993. For endorsements of (KS) see Welbourne 1979, 1983, 1994, Coady 1992, Evans 1982, and Williamson 1996. (KN) as stated requires qualification to adequately capture the spirit of what is intended. For instance, it should be qualified to take into consideration cases where the speaker is just passing on what someone else has said, where the other person knows that P but the speaker does not, perhaps because the speaker does not understand or believe the message that he is passing on. Children are often in this position, repeating to others what their parents have told them, and perhaps people hypnotized to say something when prompted. But since I am not concerned in the present paper to argue against (KN), I will pass over these matters here. For discussion, see my forthcoming-b. Three possible qualifications to (KS) will be considered in what follows, two in §IV (one having to do the speaker, another with the hearer) and one in note 26.

²See Dretske 1969, 1971, 1981.

³The locution 'presentation-as-true' is due to Burge (1993), and is meant to cover what is *implied* as well as what is *said* when a speaker makes an assertive utterance.

⁴Controversy surrounds what it is to *understand* a speaker as having made an assertion or presentation-as-true and what it is to form a belief on the basis of so understanding. Concerning understanding, one position holds that it is a form of *belief*, a belief to the effect that the speaker

asserted or implied that P (Jack 1994). Another position holds that understanding is a form of *perception* (McDowell 1979, Fricker 1987). A third position holds that understanding is a distinct cognitive state or *propositional attitude* (Burge 1993, 1997). Although I find the third answer attractive, the second is compatible with what I argue for here, and the first, though taken one way is not within the spirit of what I argue for here (nor within the spirit of the view I argue against), is not strictly incompatible with the view I advance. Deciding between these three positions, fortunately, will not matter for present purposes, for neither the truth of the standard view nor of the proposed alternative turns on any of these three answers. For some, albeit brief, discussion of what it is to *accept* a speaker's statement, see my forthcoming-b.

⁵See my forthcoming-b.

⁶See Carnap 1967, §49 on "scientific indicators", Grice 1957 on "natural meaning", and Dretske 1981, 1991.

⁷See Dretske 1981 and Goldman 1976.

⁸What makes an alternative possibility relevant? This is a difficult question that I cannot address here in any depth. When it comes to particular cases discussed in the text what matters is whether there is some other possible cause of the subject's understanding a speaker as having asserted that P other than (in the final explanation) the fact that P, some other factor which would bring it about that the hearer's basis for believing that P would fail to track the truth or co-vary with the facts. Something is a relevant alternative when, in the circumstances, it makes it likely that the grounds would be as they are (of a certain type) even though things would not be as the grounds "say" they are. Compare beliefs being of the same type. If we were, *per impossible*, to think of an experience as a proto-belief, then something would be a relevant alternative when, in the circumstances, it makes it likely that the subject would proto-believe that P even if not P. What is difficult is accounting for what makes it the case that two experiences, two so-called "proto-beliefs," are of the same type. I do not here offer a general account of perceptual or cognitive equivalence. I hope the remarks in this note and elsewhere in the text are suggestive. See also notes 14, 15, and 19. For some related discussion, see Dretske 1991. Compare Lewis 1996.

⁹Why not hold, with Nozick 1981, that the belief, and not just the grounds for the belief, must "track the truth" or carry the information that P? The second counter-example to (KS) shows why this requirement would be too strong. Someone can know that P even if his or her belief does not "co-vary" with the truth in the right way; only the grounds need so co-vary. See notes 18 and 20.

¹⁰For some discussion, see my forthcoming-a. See also Fricker 1987, 1994, Peacocke 1986, Burge 1993, McDowell 1994, Foley 1994, and Audi 1997.

¹¹Is the idiom of information carrying essential to the making of my case? I think the answer is no. What is essential is that the believer's grounds guarantee that P or establish the fact that P. One can give an account of such grounds without using the information-carrying idiom. One can stick with talk of "truth-tracking" (Nozick 1981) or "reliability" and "no relevant alternatives" (Goldman 1986) or "proper functioning in normal global and local contexts" (Plantinga 1993, 1997), or in some other truth-conducive terms. Or one could just talk in terms of guarantees without accounting for why the guarantees obtain (McDowell 1994, Fogelin 1994). One can use these notions and still end up with an alternative to the standard view, for there will be cases where speakers know that P but fail to guarantee that P when stating that P, and cases where the speaker does not know that P but nonetheless guarantees that P when stating that P. I will return to this point below in §VI. However, I am convinced that once one goes deeper and tries to account for the guarantees, one is forced to turn to the laws operative *in situ*, and one ends up endorsing the information-carrying view, whether one uses the words 'information carrying' or not. Furthermore, I find the idiom useful for it is seamless with theories of content and also allows for simple descriptions of the examples discussed here.

¹²But see also footnote 20 that contains a third *kind* of counter-example.

¹³In Harman's version, the version Adler discusses, the other reporters report what is *not* the case. Although the example as presented by Harman and Adler would serve to make my point, I have altered it so that it is structurally analogous to the barns case in a certain respect, and so that it is very clearly analogous to the thermometer case from Goldman 1986, discussed below in note 15.

¹⁴There is another way of describing the case where she *does* come to know by relying on *Andy*. Suppose Jenny *only* reads Andy's paper. Suppose she has made a commitment to him and his paper. Perhaps it is the paper her family has read for years or perhaps she has decided to trust only Andy. If his paper were not available, she would not read any other, or if she did see another paper, she would remain agnostic about its contents. The other papers are not cognitively equivalent for her. They are not relevant alternatives for her because they would not put her in a state of understanding *Andy* as having reported that P. In that case it does seem that she comes to know by relying on Andy. It is as if the other newspapers do not exist; they have no effect on her.

¹⁵The case is structurally the same as a case due to Goldman (1986: 45–6). Call it the *Thermometer Case*. Suppose a father reaches into a box of thermometers and picks the only one that happens to work; all the rest would read 98°F regardless. Although his child's temperature is 98°F and the father believes it is so on the basis of using the reliable thermometer, he does not know that his child's temperature is 98°F. This case, like Jenny's case, might be thought of as a counter-example to information-carrying theories generally. Indeed, such was Goldman's (and Fogelin's 1994) intent when offering the example. But this is not so. The mistake is confusing the state of the thermometer, the level of the mercury, with the basis for the father's belief. The basis, strictly speaking, is *not* the state of the thermometer, but the father's internal state of seeing or registering that the thermometer reads 98°F. What the information-theoretic principles require is that the believer form his or her belief on the basis of an internal ground that carries the information that P. Since the father would just as likely have picked up an unreliable though accidentally accurate thermometer, his perceptual state does not carry the information that the child's temperature is 98°F. The thermometer he uses may "track the truth" though his perceptual state does not.

On the other hand, suppose that the reliable thermometer is the only one with a purple label, and suppose the father would only use that one. He would never use any of the others. For him they might as well be swizzle sticks for cocktails. He would then be like Jenny in the previous note who only reads Andy's column. Then it seems the father would know when using the reliable thermometer. Thanks to Fred Dretske for this last point.

¹⁶I have varied Dretske's example by adding that Chianti was also served at the party, hence it should be clear that in this context the possibility that George was served a Chianti, and not a Bordeaux, is relevant.

¹⁷Wine aficionados inform me that it is difficult to imagine how George could come to mistakenly believe that Tuscany is in Bordeaux given the extensive labeling on bottles of wine. This, I believe, is an accidental feature of the example. We could easily imagine a system of labeling that did not make it hard to imagine someone with George's discriminating palate falsely believing that Tuscany is in Bordeaux.

¹⁸(1992: ch. 13). Coady argues that there is no *asymmetry* between George and Michael. He argues first that *neither* know and then that *both* know. Since the underlying reason is the same in both cases, I here only discuss his argument that neither know. Coady gives an additional reason. He says that Dretske is committed to Nozick's condition (3), the truth-tracking condition, on knowledge. Given this, it follows that George does not know, hence neither knows. The reason George does not know is that he would believe that it was a Bordeaux even if it were a Chianti. Hence his belief does not "track the truth" (1992: 225ff.). Coady's reasoning is correct; the example is, as I see it, a counter-example to Nozick. The attribution, however, is incorrect, for Dretske holds that whereas the perceptual state must "track the truth" or carry the information, the belief that results need not. In effect Dretske holds that Nozick's tracking condition (3) is too strong on belief, but not on evidence (Dretske 1981: 90–1). Adler (1996: 105) seems to make this same interpretive mistake. Adler's mistake is egregious, for he endorses Dretske's intuitions about the example but then goes on to use Nozick's (3) to explain why Michael does not know.

¹⁹However, Coady is putting pressure on the right spot. To show that George does not know, one must show that the presence of Chianti is a relevant alternative for George, an alternative that his evidence does not rule out. To do this, Coady must show that the two taste-sensations, though different in certain respects, are perceptually equivalent for George in the way that a small barn and a large

façade might be perceptually equivalent for the driver in the countryside. Though I have not here offered a theory that would allow us to draw the line between George and the driver such that the two taste sensations are not equivalent but the two barn appearances are, it seems that such a result is clearly the right result. The following remark may help. For George, the two sensations play different *cognitive roles*. One disposes him to believe he is tasting Médoc, the other to believe he is tasting Chianti. The former may lead him to order more if he is eating one kind of food, and the latter may lead him to drink more water or order Italian. All of this is *evidence* that the two sensations are not perceptually equivalent for George.

²⁰There are other kinds of cases where the speaker will know that P but will not guarantee that P when stating that P because his belief that P will not carry the information that P. In the example in the text, this happens because of the speaker's false connecting belief. But someone can also use an *unreliable method*, as well as a reliable one, to come to believe whether P, and so sometimes know that P even though his or her belief does not itself guarantee that P. The unreliable method is like the false connecting belief. Here is an example from Peacocke (1986). Mary sometimes believes that it is raining because she looks and sometimes because she consults an astrological table. When she looks outside and sees that it is raining and tells you over the phone that it is raining, you don't come to know that it is raining even though she knows. She would tell you that it is raining even if it were not. What Mary has are two distinct kinds of mental states, one a perceptual experience of the weather and the other a cognitive state of understanding an astrological chart, one that guarantees that it is raining and the other that does not. Because she sometimes relies on the latter, her belief that it is raining will not carry the information that it is raining, even though seeing that it is raining will. Unreliable methods, as well as false connecting beliefs, can generate counter-examples to (KS) that support (IN).

²¹Coady seems ambivalent between holding that Dretske is wrong and (KS) is fine as it stands and holding that Dretske is correct but (KS*) can treat the example. Further, since Coady holds that "competency" is a condition upon testifying, something like (KS) and (KS*) may be equivalent for Coady. "Any plausible" transmission principle, Coady writes, "should refer not only to sincerity but to the speaker's competence as a witness... Testifying is just one speech act amongst many which has a competency condition" (1992: 229). That is, a speaker S must possess "the relevant competence, authority, or credentials to state truly that P" in order to count as testifying that P (1992: 42). I have argued elsewhere (1997; cf. Fricker (1995: 396–7)) that this is false, that competence is not a necessary condition on the speech act of testimony.

²²These considerations also show that George is "trustworthy" about European wines, "globally reliable" about European wines, and so on. Lehrer's (1987, 1990, 1994) account, which holds that the hearer must know or justifiably believe that the speaker is "trustworthy," faces the same dilemma.

²³'Competence' is certainly not normally understood in this way, as the remarks in the preceding paragraph show. Compare Fricker (1987, 1994) and Cooper (1987: 85–7).

²⁴I am grateful to Peter Kung for this point.

²⁵For further discussion of this point, see my forthcoming-b.

²⁶There is one last proposal I should mention in a footnote. Instead of adding conditions on the speaker or the hearer, one might simply add to (KS) the condition that there be no defeating conditions. For instance, Evans endorses (KS**).

(KS**) If the speaker S has knowledge of x to the effect that it is F, and in consequence utters a sentence in which he refers to x, and says of it that it is F, and if his audience A hears and understands the utterance, and accepts it as true (and there are no defeating conditions), then A himself thereby comes to know of x that it is F. (Evans 1982: 310–1)

This proposal faces a similar dilemma as Coady's proposal. Either the "no defeating clause" is met but the counter-examples still go through, or the counter-examples do not go through but the clause amounts to nothing more than adding the information-carrying requirement to the antecedent of (KS). I suppose the latter is the case, for the defeating conditions are equivalent to the presence of relevant alternatives: the Chianti in George's case and the fake reports in Andy's case.

²⁷See Welbourne 1983, 1986: 6–7.

²⁸What about the fact that the boy has so often lied? Doesn't this comprise "misleading evidence" that the stranger "does not possess" that would defeat his knowledge (cf. Harman 1973)? I agree with Welbourne (1986: 7) that the answer to this question is no. It does not matter to whether the *stranger* knows, though *perhaps* it would matter to a villager who accepted the boy's present report.

²⁹In his *Inquiry into the Human Mind*, chapter VI, section XXIV, he writes: "Now, if we compare the general principles of our constitution, which fit us for receiving information from our fellow-creatures by language, with the general principles which fit us for acquiring the perception of things by our senses, we shall find them to be very similar in their nature and manner of operation."

³⁰I am grateful to Marshall Swain, Peter Kung, and one of the referees for *Noûs* for pressing this point.

³¹"The wise and beneficent Author of Nature, who intended...that we should receive the greatest and most important part of our knowledge by the information of others, hath, for these purposes, implanted in our natures two principles that tally with each other.

The first of these principles is, a propensity to speak truth... This principle has a powerful operation, even in the greatest liars; for where they lie once, they speak truth a hundred times. (...) Lying, on the contrary, is doing violence to our nature... Speaking truth is like using our natural food...but lying is like taking physic. (...) [This] original principle implanted in us...may be called the *principle of veracity*. (...) [It] is the counterpart to the...*principle of credulity*" (*Inquiry*, chapter VI, section XXIV).

³²"The internal principles and motives may operate in a uniform manner...; in the same manner as the winds, rain, clouds, and other variations of the weather are supposed to be governed by steady principles; though not easily discoverable by human sagacity and enquiry" (*Enquiry Concerning Human Understanding*, VIII, I, 68).

³³"It seems almost impossible, therefore, to engage either in science or in action of any kind without acknowledging the doctrine of necessity, and this inference from motives to voluntary actions, from characters to conduct" (*Enquiry*, VIII, I, 70).

³⁴Much of the hostility to the claim that testimony is much like perception from the epistemic point of view is driven, I think, from the justified belief that testimony can, and often does, go wrong in a number of ways. People lie, people talk about things they know little about, they are often biased, and so on. Furthermore, just as you may misspeak, I may misunderstand what you have said. The chain of communication is another source of possible error. Though these facts are true, it does not show that testimony is different in *kind* as opposed to *degree*, from the epistemological point of view, from perception. Perception, after all, is also susceptible to numerous kinds of errors or malfunctions. I hope to pursue this point, and its connection to debates surrounding when we are entitled to accept another person's report, elsewhere.

³⁵Audi (1997), who endorses (KN) at least seven times, endorses (KS) only once (p. 412), and when he does he enters the qualification "normally", with reference in a footnote to Dretkse's counterexample (1982) and Coady's (1992) rejoinder. Audi, then, seems aware that (KS), strictly speaking, is false as an explanatory principle. Burge (1993, p. 486 n.24) also shows awareness of the need to enter qualifications to principles like (KN) and (KS) to deal with the kinds of cases raised here.

³⁶It may be suggested that the standard view enjoys some explanatory power on the grounds that the standard view, though perhaps not strictly speaking true, explains why people trust speakers, for when they trust a speaker they believe that the speaker knows what it is they come to believe. Though this may be true in a number of cases, I do not think it is what explains why we believe what other people tell us. It certainly does not explain why children, who do not have the concept of knowledge, believe what their parents tell them. Why do we accept what we are told? This is a difficult and complicated question that I hope to address at length elsewhere. My hunch is that an adequate answer will incorporate aspects of the information-theoretic approach generally.

³⁷I am grateful to Kent Bach, Tyler Burge, Liz Camp, Tony Coady, Ben Escoto, Tony Genova, Peter Godfrey-Smith, Ara Lovitt, Simon May, John Perry, Debra Satz, and Ken Taylor for very helpful comments. I am most grateful to Fred Dretske and Peter Kung for numerous helpful comments on

previous drafts. Some of the material contained here was presented to audiences at Stanford University, the University of Kansas and St. Louis University. I am grateful for the discussions that followed both presentations. An earlier version of this paper was presented at the April 1998 Central Division Meeting of the APA in Chicago. I am grateful to the members of the audience, especially Marshall Swain, for their comments. I am also indebted to the two referees for *Noûs* for comments that led to substantial improvements.

References

- Adler, Jonathan (1996) "Transmitting Knowledge" *Noûs*, vol. 30, no. 1: 99–111.
- Armstrong, David M. (1973) *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Audi, Robert (1997) "The Place of Testimony in the Fabric of Justification and Knowledge" *American Philosophical Quarterly*, vol. 34, no. 4: 405–422.
- Burge, Tyler (1993) "Content Preservation" *Philosophical Review*, vol. 102, no. 4: 457–488.
- Carnap, Rudolf (1967) *The Logical Structure of the World*. Berkeley: University of California Press.
- Coady, C. A. J. (1992) *Testimony: A Philosophical Study*. Oxford: The Clarendon Press.
- Cooper, David (1987) "Assertion, Phenomenology, and Essence" *Proceedings of the Aristotelian Society*, Supplementary Volume 61: 85–106.
- Dretske, Fred (1969) *Seeing and Knowing*. New York: Routledge, Kegan & Paul.
- Dretske, Fred (1971) "Conclusive Reasons" *Australasian Journal of Philosophy*.
- Dretske, Fred (1981) *Knowledge and the Flow of Information*. Cambridge, Mass.: MIT Press.
- Dretske, Fred (1982) "A Cognitive Cul-de-sac" *Mind*, vol. 91, no. 361: 109–111.
- Dretske, Fred (1991) "Knowledge: Sanford and Cohen" in *Dretske and His Critics*, McGlaughlin ed., Oxford: Basil Blackwell Publishers.
- Dummett, Michael (1994) "Testimony and Memory" in Matilal and Chakrabarti (1994).
- Evans, Gareth (1982) *The Varieties of Reference*. Oxford: Oxford University Press.
- Fogelin, Robert (1996) *Phyrronian Reflections on Knowledge and Justification*. New York: Oxford University Press.
- Foley, Richard (1994) "Egoism in Epistemology" in Frederick Schmitt, ed., *Socializing Epistemology: The Social Dimensions of Knowledge*. Rowman and Littlefield: 53–73.
- Fricker, Elizabeth (1987) "The Epistemology of Testimony" *Proceedings of the Aristotelian Society*, Supplementary Volume 61: 57–84.
- Fricker, Elizabeth (1994) "Against Gullibility" in Matilal and Chakrabarti (1994).
- Fricker, Elizabeth (1995) "Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony" *Mind*, vol. 104, n. 414.
- Goldman, Alvin (1976) "Discrimination and Perceptual Knowledge" *Journal of Philosophy*, vol. 73, no. 20: 771–791.
- Goldman, Alvin (1986) *Epistemology and Cognition*. Cambridge, Mass.: Harvard University Press.
- Graham, Peter J. (1997) "What is Testimony?" *Philosophical Quarterly*, vol. 47, no. 187.
- Graham, Peter J. (forthcoming-a) "The Reliability of Testimony" *Philosophy and Phenomenological Research*.
- Graham, Peter J. (forthcoming-b) "Conveying Information" *Synthese*.
- Grice, H. P. (1957) "Meaning" *Philosophical Review*, vol. 66.
- Hardwig, Jonathan (1985) "Epistemic Dependence" *Journal of Philosophy*, vol. 82, no. 7.
- Hardwig, Jonathan (1991) "The Role of Trust in Knowledge" *Journal of Philosophy*, vol. 88, no. 12.
- Harman, Gilbert (1973) *Thought*. Princeton, N.J.: Princeton University Press.
- Jack, Julie (1994) "The Role of Comprehension" in Matilal and Chakrabarti (1994).
- Lehrer, Keith (1987) "Personal and Social Knowledge" *Synthese*, vol. 73, no. 1.
- Lehrer, Keith (1990) *Theory of Knowledge*. Boulder: Westview Press.
- Lehrer, Keith (1994) "Testimony, Justification and Coherence" in Matilal and Chakrabarti (1994).
- Lewis, David (1996) "Elusive Knowledge" *Australasian Journal of Philosophy*, vol. 74, no. 4.

- McDowell, John (1979) "Anti-Realism and the Epistemology of Understanding" in Margalit, ed., *Meaning and Use*.
- McDowell, John (1980) "Meaning, Communication, and Knowledge" in Zak van Straaten, ed., *Philosophical Subjects*. Oxford: The Clarendon Press.
- McDowell, John (1994) "Knowledge by Hearsay" in Matilal and Chakrabarti (1994).
- Matilal, B.K. and A. Chakrabarti (1994), eds., *Knowing From Words*. Dordrecht: Kluwer Academic Publishers.
- Nozick, R. (1981) *Philosophical Explanations*. Cambridge, Mass.: Harvard University Press.
- Peacocke, Christopher (1986) *Thoughts: An Essay on Content*. Oxford: Basil Blackwell.
- Plantinga, A. Ivin (1993) *Warrant and Proper Function*. New York: Oxford University Press.
- Plantinga, Alvin (1997) "Warrant and Accidentally True Belief" *Analysis*, vol. 57, no. 2.
- Ross, Angus (1986) "Why Believe What We Are Told?" *Ratio*, vol. 28, no. 1.
- Ross, James (1975) "Testimonial Knowledge" in *Analysis and Metaphysics*, Lehrer ed., Dordrecht: Reidel.
- Swain, Marshall (1981) *Reasons and Knowledge*. Ithaca: Cornell University Press.
- Welbourne, Michael (1979) "The Transmission of Knowledge" *The Philosophical Quarterly*, vol. 29.
- Welbourne, Michael (1983) "A Cognitive Thoroughfare" *Mind*, vol. 92, no. 367: 410–12.
- Welbourne, Michael (1986) *The Community of Knowledge*. Aberdeen.
- Welbourne (1994) "Testimony, Knowledge, and Belief" in Matilal and Chakrabarti (1994).
- Williamson, Timothy (1996) "Knowing and Asserting" *The Philosophical Review*.