



The Epistemology of Testimony

*Peter Lipton**

1. Introduction

Is there anything you know entirely off your own bat? Your knowledge depends pervasively on the word of others. Knowledge of events before you were born or outside your immediate neighborhood are the obvious cases, but your epistemic dependence on testimony goes far deeper than this. Mundane beliefs—such as that the earth is round or that you think with your brain—almost invariably depend on testimony, and even quite personal facts—such as your birthday or the identity of your biological parents—can only be known with the help of others. Science is no refuge from the ubiquity of testimony. At least most of the theories that a scientist accepts, she accepts because of what others say. The same goes for almost all the data, since she didn't perform those experiments herself. Even in those experiments she did perform, she relied on testimony hand over fist: just think of all those labels on the chemicals. Even her personal observations may have depended on testimony, if observation is theory-laden, since those theories with which it is laden were themselves accepted on testimony. Even if observation were not theory-laden, the testimony-ladenness of knowledge should be beyond dispute. We live in a sea of assertions and little if any of our knowledge would exist without it.

If the role of testimony in knowledge is so vast, why is its role in the history of epistemology so slight? Why doesn't the philosophical canon sparkle with titles such as *Meditations on Testimony*, *A Treatise Concerning Human Testimony*, and *Language, Truth and Testimony*? The answer is unclear. One might try to explain the neglect of testimony by appeal to an individualistic bias in the history of western epistemology, but this bias, though real, is not very explanatory, since 'epistemic individualism' is little more than another expression for the neglect of testi-

* Department of History and Philosophy of Science, University of Cambridge, Free School Lane, Cambridge CB2 3RH, U.K.

Received 15 November 1996; in revised form 21 July 1997

mony. A more satisfying explanation might be that testimony has been seen as derivative, a means merely of the transmission and not the creation of knowledge, and so a factor that fundamental epistemology can safely ignore.

This may help to explain the neglect of testimony, but it does not justify that neglect. The transmission of knowledge is or ought to be of central epistemological interest. Knowledge that depends on memory, for example, might profitably be viewed as the transmission of knowledge from past to present self, but this would hardly suggest that the epistemology of memory is unimportant. In any case, testimony is a means of the creation of knowledge. From an individualistic perspective this is obvious, since what I learn from others is new knowledge for me; but it is almost as obvious from a social perspective, as one can see by trying to imagine science without communication. Testimony ought to be one of the central topics in the theory of knowledge.

Fortunately, the neglect of testimony may have come to an end, thanks in large measure to two recent books on the subject: Tony Coady's *Testimony* (1992) and Steven Shapin's *A Social History of Truth* (1994). Both provide sustained studies of the epistemology of testimony; both are fine pieces of work. Neither book, it must be said, provides ideal bedtime reading. Both authors like to circle their quarry, returning to the same issues again and again from various angles: their discussions are rich, dense and rewarding. My aim in this article is to develop some of my own thoughts about the epistemology of testimony by grappling with theirs.

The central theme that Coady and Shapin share is the one I have already endorsed: the ubiquity of testimony. A signal virtue of both books is the way they illuminate the enormous extent to which we rely on each other's word, the extent to which the acquisition of knowledge is a collective enterprise. This theme is not just a common starting point, but a common motif that runs the length of both books. What the two authors make of this motif, however, is interestingly different. Coady mainly uses the ubiquity of testimony in two ways. The first is to argue against the possibility of any reductive account of the epistemology of testimony that would justify our reliance on testimony in terms that are themselves testimony-free. David Hume's attempt to reduce testimonial inference to simple enumerative induction in his discussion of miracles provides Coady's most important target. The upshot of Coady's attack on reductionism is a kind of fundamentalism about testimony, according to which testimony is an irreducible source of knowledge, in many respects on a par with perception, memory, and deductive inference. The second main use Coady makes of the ubiquity of testimony is to explore the prospects for non-reductive vindication of our reliance on testimony. The emphasis here is on a style of transcendental argument, which aims to show that testimony must be reliable if a public language is possible at all. These two strands of Coady's discussion seem to me the most central and important in his book, but it contains much more. Virtually every philosopher in the standard philosophical canon who has written on testimony is discussed, and there are chapters on empirical work

on the reliability of testimony, the role of expert witnesses in the law, and much more besides.

For Coady, the ubiquity of testimony raises the question of justification; for Shapin it raises the question of management. Given the extent to which science depends upon testimony, the crucial question for Shapin becomes: which testimony should one accept? This is a question that every scientific community must answer, but different communities may answer it in different ways. The community that Shapin analyses in depth is that of the Royal Society of seventeenth-century England, with Robert Boyle in the starring role. Shapin argues that the decision on what to believe is largely a question of what kind of person to trust and, in the community of the Royal Society, the answer to that question was, broadly speaking: a gentleman. Gentlemen were counted trustworthy because of their moral qualities. They were by their nature counted free and virtuous: free men had no motive to lie, and virtuous men were committed to truthfulness. Moreover, because of the honour code that governed gentlemanly society, the cost of not trusting another gentleman was impossibly high. Here too, the central themes of the book are embedded in a rich discussion that raises many other important issues. Shapin has a great deal to say about what made for gentlemanly status, about Boyle's career and the way that career was presented, about the handling of awkward testimony in the cases of hydrostatics and astronomy, about the practical and epistemic role of servants and laboratory assistants, and about much else.

In part because of Coady and Shapin's fundamental agreement over the ubiquity of testimony, and their consequent agreement about the need to reject individualistic epistemologies, the contrasts between the two authors in both approach and doctrine are particularly interesting. Shapin is a social constructivist and a relativist, at least on matters of method; Coady has no patience with relativism, though he is no naive realist. Coady emphasises the lack of criteria for accepting testimony and the universality of a policy of largely indiscriminating trust; Shapin is especially interested in how discriminations in trust are made, and in the particularities of the way the universal need to discriminate is met differently by different communities. Coady's epistemological concerns are largely normative; Shapin, it seems to me, attempts a kind of reduction of epistemology to descriptive ethics. These are deep contrasts, but a number of them are complementary so that, read in tandem, these books often seem to fill out different parts of a large and largely shared picture, rather than to offer incompatible representations of the role of testimony in our lives.

Where the two books are in tension, I found myself siding sometimes with one, sometimes with the other, though my own work and approach is much closer to Coady's than to Shapin's. It seems to me that Shapin doesn't care enough about truth and Coady doesn't care enough about the importance of discriminating reliable from unreliable testimony. Shapin doesn't put sufficient emphasis on the testifier's competence; Coady doesn't worry enough about the testifier's honesty.

Coady's epistemology would make us too gullible; Shapin's too selective. Moreover, although neither author rules out a 'cognitive' or inferential account of the epistemology of testimony, both occlude such an account, Coady by emphasising simple acceptance of what we are told, Shapin by the reduction of epistemology to ethics. I try to defend some of these claims in what follows, beginning with some critical remarks about Shapin's book, then about Coady's, and ending with some reasons why we might want a cognitive solution to the problem of deciding whom to believe and a suggestion about where we might look for one.

2. Truth and Method

To give striking support to an approach or a theory, in any area of inquiry, take a phenomenon that the approach seems especially ill suited to handle and show how the phenomenon can be reconfigured so as to provide a successful application of that approach. For example, altruistic behaviour is on the face of it a phenomenon that evolutionary biology is particularly ill suited to explain, and this is one reason for the fascination with sociobiological accounts, which attempt to show that various sorts of altruism are precisely what an evolutionary mechanism should be expected to produce. What is attempted in cases like this is the conversion of an apparent counterexample into confirmatory evidence.

Shapin's project can be seen as a kind of counterexample conversion. However plausible a social constructivist story of the negotiations that lead to the acceptance of this or that bit of high theory, the constructivist approach seems ill suited to account for the role of relatively mundane observation. The example of the Royal Society, with its emphasis on individualist empiricism, might be thought a particularly awkward case for the claimed ubiquity of testimony. Their motto, after all, was *Nullius in verba* (On no man's word). One of Shapin's central aims is to show that dependence on testimony was ubiquitous in this case, as in all others, that it hence required extensive management, and that the social constructivist approach is well suited to illuminate how this management was achieved. The extent to which Shapin succeeds in this is impressive. One reason it is so successful is that Shapin is especially good at balancing the particular against the universal, at revealing the particular and contingent ways that one scientific community handles the universal problem of the management of testimony.

For all my admiration for the historical achievements of social constructivists, however, I disdain their relativism, and I will begin my highly selective critique of *A Social History of Truth* with that topic. Shapin presents his book from the point of view of what he calls a 'liberal' notion of truth, a notion that equates truth and consensual belief. He calls the notion 'liberal' because it allows that what is true will vary between times and between communities, since consensual belief so varies. He contrasts this relativism with what he call the 'restrictive' notion of truth, which distinguishes between truth and belief, and maintains that what is true does not vary over time or space. The restrictive notion is naturally associated with

correspondence theories of truth, which makes what is true depend on the state of a mind-independent reality, but it would also cover some other standard philosophical accounts, including various Kantian, redundancy, and pragmatic theories. There are also members of the philosophical menagerie that are neither 'liberal' nor 'restrictive', such as some coherence, instrumental, and again some pragmatic theories. There is logical space for these theories of truth that are neither liberal nor restrictive because one can deny the equation between truth and consensual belief without maintaining that the truth is the same for all communities.

The *frisson* that the title 'A Social History of Truth' is intended to provoke depends upon the contrast between the 'liberal' perspective Shapin adopts and the 'restrictive' perspective that he takes it many of his readers will have (though committed social constructivists will presumably take the title to suggest a truism). Partly for this reason, and partly because I am into epistemic bondage, I want to make some comment here on this very large topic. Because Shapin's taxonomy is not exhaustive, and because I find his terminology uncongenially loaded, from here on the 'liberal' view will be referred to as 'relativism', and any of the other views as 'realism'.

The first question to ask is whether Shapin's relativism matters to his project. Shapin's own answer is no and yes. He claims that nothing in his book provides an argument against realism (Shapin, 1994, p. 4), and I agree. Realists will need to translate the title as *A Social History of Belief*, which will eliminate the *frisson*, but will not I think alter the substantive message. Similar translations will be necessary from time to time in the body of the book, but they strike me as equally innocuous and only very seldom required. Certainly, a realist may read the book with great profit and may even accept almost all of its claims without metaphysical qualm. This is just what one would expect: the relativist and the realist both deploy the notion of belief, only the realist entertains a distinct notion of truth that the relativist rejects. So if the relativist tells a story about the dynamics of belief, there need be nothing in the story that the realist must reject.

Nevertheless, Shapin also maintains that his relativism matters methodologically, even if it makes little substantive difference to his argument (Shapin, 1994, pp. 4–5). He claims it matters because realism creates a bias in favour of our own stock of knowledge, blocking curiosity about cultural variation. Relativism, by contrast, disposes us towards a charitable reading of alien belief and encourages us to leave aside what we believe the world to be like when trying to understand what historical actors thought. It also, according to Shapin, encourages the historian to adopt a salutary historicism, which interprets historical actors in their own terms (p. xvi). Shapin seems to acknowledge that, as ordinary actors, we are all realists, but he advocates relativism as an important resource for the special project of interpreting cultural variation. His willingness to tolerate both relativist and realist notions is I think also meant to be the consistent application of his relativism one level up.

Shapin's case for relativism in this book is very briefly made, but as given I find the arguments unpersuasive and the position unattractive, though I too will only offer a brief case. Since the notion of belief is common property between realists and relativists, realists are perfectly free to investigate the history of belief and its variation. Nor is it at all clear why a belief that there is only one truth should make one less interested in the question of why different people have believed different things. Moreover, insofar as a relativist can explain the development of past belief without appealing to what she knows but her subjects did not, so can the realist.

To this last claim the relativist may respond that, even if this feat of epistemic self-control were possible in principle, the realist would not in practice manage to avoid any appeal in his historical explanations to how he, deep in his heart, knows the world to be. Perhaps this would be difficult, but I do not see why it should be any more difficult than the feat that the relativist needs to perform, of suspending her own realism as an ordinary actor when she turns her hand to history. (Maybe both are impossible.) Indeed it is not even clear that the two feats are distinct.

Nor would it help the relativist if she could establish or persuade herself that relativism just is the true account of truth. Paradox aside, this will not serve the purpose, because epistemic relativism does not entail epistemic tolerance, any more than moral relativism entails moral tolerance. I may accept that what you believe to be morally right is indeed right for you even if it differs from what I take to be right, yet decide that one of the things that is right for me is ruthlessly to oppress you. I may decide that what your community believes is indeed true for you, even though it contradicts what I believe, yet decide to explain your beliefs in my terms and by appeal to how I take the world to be. The case for historicism must be made apart from the case for relativism, and if it can be so made then it can be accepted by the realist.

This suggests why I hold that the realist is in no worse shape than the relativist, when it comes to historical methodology. But there are also special liabilities that the relativist faces. At first glance, a relativist view of truth seems strikingly inappropriate to Shapin's particular project, which is to explain how a particular community managed the problem of determining which testimony to accept and which to reject, because that choice presupposes a distinction between what is believed (or at least asserted) and what is the case. This point is related to one of Plato's central objections to relativism in the *Theaetetus*, his argument from judgements about the future (171d–172c, 177c–179b). Even if relativism were acceptable for judgements about current appearances, so that for them what seems true is the same as what is true, it would not be acceptable for judgements about the future, even judgements about future experience. For in this case, the judger cannot contain the criterion of truth within herself: she must admit the possibility of a mismatch between what she believes she will experience and the experience when it comes. Similarly, when we evaluate testimony, we must acknowledge the possibility of

justified false belief: the possibility of the gap between what a correct application of our epistemic standards might lead us to accept and what we might later discover. Thus those who rely on testimony but are not completely gullible must make the realist's distinction between belief and truth. So it may seem that one cannot make sense of the selective use of testimony from a relativist's point of view. But this is too quick. As we have seen, Shapin is careful to distinguish between the attitude of the actor and the special-purpose attitude of the analyst. What I have just said shows that choosy consumers of testimony must be realists, but this does not immediately preclude selective relativism for historical purposes.

That special-purpose relativism is nevertheless problematic. First of all, it seems to violate the principle of historicism, since it deploys a relativist's notion of truth that is alien to the historical actors. The problem is not just that those actors did not in fact deploy this notion: no actors can consistently think relativistically about their own ordinary judgements and discriminations. Secondly, relativism generates an inconsistency that Shapin's 'special-purpose' locution elides. As an historian who relies hand-over-fist on testimony but is at least as choosy as most, Shapin must be a realist about his own claims, making a sharp distinction between what members of his community may generally hold and what is really the case about past actors. At the same time, his methodology instructs him to deny the distinction. The only way I can see to avoid the inconsistency is to maintain the distinction between truth and belief, while attempting a principled refusal to explain the behaviour of past actors in terms they did not deploy. And that is simply the realist-historicist's approach to the problem.

3. The Morality of Epistemology

We turn from the true to the good. Shapin encourages us to see the epistemology of testimony as a problem in applied ethics. He argues for this at two levels, the global and the local. The global arguments are that testimony is by its nature a moral issue, since accepting testimony is a matter of trust, since the social order depends upon it and since giving testimony is tantamount to making a promise that one's word can be relied upon. Testimony inherits the rich ethical content of placing trust, maintaining the social order and keeping promises. The local arguments concern the criteria of credibility that the seventeenth century gentlemen deployed. These criteria were moral, based as they were on the concepts of free action, honor, virtue, and civil conversation.

Shapin is very successful at eliciting moral aspects of testimony, and this is from my point of view a particularly important achievement, since the connection between epistemology and ethics is almost entirely ignored by the standard epistemological literature. Shapin's discussion and references also suggest that Coady would have found the philosophical literature on testimony somewhat less sparse had he investigated sources in ethics and political theory. At the same time, while Shapin places the existence of a substantial moral component in the epistemology

of testimony beyond doubt, his discussion also suggests the radical programme of reducing the epistemology of testimony reduces to ethics. I want briefly to consider this reductive idea.

A claim that the epistemology of testimony reduces to ethics can mean various things. One is that all the criteria for the credibility of testimony are overtly criteria of moral assessment. Shapin's specific arguments concerning seventeenth-century gentlemen's criteria of credibility and the moral elements they involved would not support this sort of general reductive claim, since only some of the criteria in play were overtly moral. To determine whether to accept a piece of testimony, the actors Shapin discusses, like the rest of us, take many factors into account. For example, we all make use of the factors mentioned in Locke's famous list (Locke, 1700, IV.XVI.4), a list that includes the integrity of the informant, a moral question, but also the number and skill of the witnesses, the 'design' of the author, the consistency of the different parts of the testimony and the existence of contrary testimonies. These, unlike honor and virtue, are not overtly moral criteria.

A second version of the reductive claim is that, while the specific criteria of credibility may be more or less overtly moral, the institution of testimony is by its nature a moral institution. This is the sort of reductive claim that Shapin's global arguments seem aimed to establish. The first universal argument is the argument from trust. The institution of testimony is by its nature moral, according to this argument, because the decision whether to accept someone's testimony is always a question of whether to *trust* the informant, and judgements of trustworthiness are always moral judgements. On its own, this argument is not persuasive. We may grant that the decision whether to accept someone's testimony can always be glossed as the decision whether to trust that person, but we do not yet have a reason to suppose that trust in this sense is always a moral issue. It seems rather to be a sense that includes scientists' decisions whether to trust their theories, instrumentation and data.

The second global argument is that the institution of testimony is by its nature moral, because it is a prerequisite for the existence of any social order. The premise seems clearly correct: given our ubiquitous dependence on the word of others, no society could exist without testimonial practices. But, one might say, nor could a society exist in the absence of sense perception or reasoning, yet this is hardly sufficient to show that the epistemology of perception or reasoning reduces to ethics: not every prerequisite for a moral practice or concept is itself moral. If this objection to Shapin's argument is unconvincing, perhaps it is in part because the point is not just that testimony is a prerequisite for social order, but that it partially constitutes that order. Testimony is itself a social activity in a way that perception and reasoning are not. (Or so it seems. Given Coady's case for the influence of testimony on all other forms of belief acquisition, however, one might in the end want to say that all of these are to some extent social.) Nevertheless, even being partially constitutive of a social practice does not entail a strong moral reduction.

Language, or its use, is of course part of social practice, but it does not follow that the philosophy of language reduces to ethics.

This leaves us with the third and I think the best of Shapin's global arguments, the argument from promises. In outline, the argument is that promising is a moral institution, to give testimony is tantamount to promising that one's word can be relied upon, so testimony is a moral institution. This argument looks strong, because promises do have an essentially moral character, and testimony does seem to be as morally charged as promising. As it happens, Kant came very close to making Shapin's argument for him. Promising is one of the examples Kant uses to illustrate his first formulation of the Categorical Imperative, according to which we should only act on a maxim that we could will to be a universal law. The Categorical Imperative is supposed to exclude false promises, because the attempt to universalise a maxim that permitted them would destroy the institution of promising. Conveniently enough, Kant equates false promising with lying, so that his claim is effectively also that we should not lie, since the attempt to universalise a maxim that permitted lies would eliminate the practice of testimony (Kant, 1785, First Section). This certainly suggests that promising and giving testimony are both essentially moral, and for the same reason.

Nevertheless, there is a salient difference between testimony and promises. When I promise to meet you for lunch tomorrow, I vouch that I will be there, not just that I would like to be. Testimony is sometimes but not always like this. When an informant speaks, she is sometimes vouching for the truth of what she says, so that she may be held morally culpable if what she said turns out to be false, even if she believed what she said. In many other cases, however, all she is morally responsible for is her belief in what she says. Testimony may be false in two ways: through lying and through honest mistake. Lying is always a moral issue, but error is only a moral issue some of the time. If I tell you the ice is safe, I take responsibility for being right, not just for being honest, but this is only one type of case. I do not suppose that everyone with whom I disagree is for that reason morally deficient. If there is a temptation to think that one always takes moral responsibility for the truth of one's testimony, this may be due to the formal ring of the word; once reminded that we are here construing 'testimony' so broadly as to include virtually all assertion, the temptation should disappear. There thus appears to be an important disanalogy between promises and testimony: false promises always raise a moral issue, whereas false assertions do not. The most one can say is that lying is always an *option* in the case of testimony, and someone who gives testimony does vouch for his honesty, so to this extent testimony is indeed a moral issue.

In Shapin's discussion, the contrast to true testimony is almost always a lie, rather than an error (cf. Feingold, 1996, p. 138). Why is this? One possibility is that his relativist methodology makes it seem inappropriate to invoke the category of error. But it is also part of Shapin's methodology to abjure appeal to the actors'

psychological states (p. xxiii), which might seem to exclude attributions of lying too, because of the mendacious intention that this requires. Indeed, he tells us to construe his use of 'belief' behaviouristically, as 'acting as if one believed', which would seem to rule out in one go the attribution of both error and of mendacity. The solution to this puzzle is to insist on the distinction I mentioned earlier between actor and analyst. Shapin will not claim that his actors were either mistaken or dishonest, but he is free to attribute to his actors both sorts of claim.

There are more plausible explanations for the focus on lying. One is that it meshes with Shapin's detailed development of the nature of the gentleman. Freedom, virtue and honor are supposed to be proof against lying, but not against error. It is the case of lying rather than the case of error that suits the local argument for the moral status of gentlemen's testimony. Another explanation for Shapin's focus on the lie is that it is the universal possibility of mendacity rather than the universal possibility of error that supports the claim that the management of testimony always has a moral component. A final explanation for the focus on the lie, however, would be an attempt to support a reduction of epistemology to ethics in a stronger form than I have so far considered. This would be the claim not merely that the decision of what testimony to accept has a moral dimension, but that moral assessment exhausts the space. It would be the attempt to replace the cognitive by the moral; but none of Shapin's arguments supports such moral imperialism. The sensible reminder of the permanent possibility of mendacity needs to be distinguished from any suggestion that all the credibility factors that consumers of testimony use to distribute their trust are themselves morally loaded. We may be quite convinced that our informant is an honest chap, yet still wonder whether to believe what he says, and the considerations that govern our decision in such cases will be epistemic but not in any straightforward sense ethical. An epistemology of testimony should illuminate these cognitive considerations, and these considerations should not be occluded by the moral factors that accompany them.

4. Gentlemen Prefer Gentlemen

Shapin's book offers a rich and detailed discussion of the nature and status of gentlemanly culture in seventeenth century Britain and of Robert Boyle's own position in that culture. I am not in a position to assess Shapin's claims properly here, but his discussion must be essential reading for anyone concerned with this aspect of cultural history. I will, however, venture a few criticisms of some of Shapin's claims about the role of gentlemanly culture in the epistemology of testimony.

Shapin presents the central features of gentlemanly culture, especially freedom of action and the codes of honor and virtue, as peculiarly effective resources for establishing relations of trust. Yet some of these features and their consequences might be liabilities rather than advantages. Two seem particularly problematic: the

freedom of action that gentlemen enjoyed and the extraordinarily high cost that Shapin claims was associated with contradicting a gentleman's testimony.

Shapin argues that gentlemen's free status encouraged trust in their testimony, primarily because it made it possible for them to present themselves as disinterested informants self-constrained by the honor code. This is plausible, but freedom cuts both ways. It also presumably allowed and was seen to allow gentlemen to get away with false testimony in a way that the comparatively unfree could not. Who had more to lose from being judged an unreliable informant, the master or the servant? Shapin's discussion of the role that the recognition of free status played in distributions of trust seems one-sided. Similarly, the high, even mortal cost of denying a gentleman's testimony seems a mixed epistemic blessing. Shapin argues, again with plausibility and insight, that this arrangement supported a culture of civil conversation that had clear advantages for the management of testimony. But it also seems that it must have had clear disadvantages. No gentleman could in fact accept everything claimed by every other gentleman on pain of pervasive contradiction, so, if trust could not be publicly withheld, the denials must have moved underground. This would and presumably did create various difficulties for the management of testimony, alongside its civil advantages. Shapin does an excellent job of illustrating some of the ways gentlemen scientists managed to avoid public contradiction, but he does not in my view say enough about the inevitable contrast between public acquiescence and private dissent (cf. Feingold, 1996, p. 134).

To my claim that some of the features of gentlemanly status that Shapin takes to be central to gentlemen's testimonial practices would have provided liabilities as well as advantages for the management of testimony, Shapin ought I think to respond that this is so but constitutes no criticism. His point is that this was the way testimony was in fact managed by that particular group of actors, whether or not it was the easiest conceivable means of getting the job done. My objection so far would then just be the mild one that Shapin's presentation is skewed towards the advantages. But I want now also to question the claim that it was the peculiar features of gentlemanly culture that provides the primary explanation of gentlemen's distribution of trust in each other. My question is based on two very general observations, both of which Shapin seems to accept and which in any case are pretty much direct consequences of the ubiquitous reliance on testimony that applies to members of any community and which both Shapin's and Coady's books bring out so well.

Given the extent of any actor's dependence on the word of others, it seems clear that he must trust most of the testimony of members of his own community, given that these members will in almost every case provide the actor with his primary source of testimony. This trust is essential if the actor is to have a workable stock of beliefs at all and, as Shapin observes, it is encouraged by the repeated and 'face-to-face' interaction that in-group testimony allows (Shapin, 1994, pp. 414–416).

Not only are we in general more inclined to trust testimony delivered to our faces, but repeated interaction with the same individuals often provides especially good opportunities to assess the reliability of informants and so particularly strong motives for informants to act reliably. Crudely put, we are especially inclined to trust people we know, because we sometimes have especially good reasons to trust them and because sometimes we have no choice.

Faced then with the question of why gentlemen trusted each other, the obvious answer is simply that they all belonged to the same club. We could say that they trusted each other because they were all gentlemen, but this answer is ambiguous between Shapin's explanation and the simpler one I have just offered. The difference comes out when we move from group to group. According to Shapin, his gentlemen scientists were often suspicious of the testimony of 'vulgar divers', who made claims about the effect of diving on perceived pressure. Yet given the enormous need for testimony from those with whom one lives, the community of vulgar divers could not but largely trust each other, however suspicious they may have been of the testimony of gentlemen. The general explanation I have suggested does not appeal directly to freedom, honor, and virtue, which might differentiate the gentlemen from the divers. The explanation is just that gentlemen, like vulgar divers, trusted their own.

The relationship between Shapin's 'honor explanation' and my 'communal explanation' for the trust that gentlemen placed in each other is not straightforward. Logically, they do not exclude each other: they could both be true, and indeed Shapin seems to endorse something like the communal explanation at the end of his book, where he presses his universal claims about testimony and trust. But one might hold that the two explanations are nevertheless epistemic competitors. We know that the communal explanation applies to the community of gentlemen, because it applies to all communities and, insofar as we have this explanation, it undermines our reason for also endorsing the honor explanation. An analogy may clarify this point. If my computer will not turn on, I might consider two explanations, that its fuse is blown or that it isn't plugged in. Both these explanations could be true, but once I see that the plug is out I lose any evidence I might have had for the blown fuse. If we take this view of the relation between the honor and communal explanations, then we ought to conclude that Shapin has at least substantially overstated his case.

It is not, however, clear to me that we ought to take the relation between the explanations in this way. An alternative, and one that Shapin would I think favour, is that the two explanations are not even epistemic competitors. Instead, the honor explanation should be seen as the communal explanation applied to a particular case. If another crude analogy is wanted, the relationship between the communal and honor explanation is not like that between the plug and the fuse explanation, but perhaps like the relationship between a Darwinian natural selection explanation for the presence of a given trait and an explanation appealing to the mechanism

of gene selection and reproduction. In this case, the second explanation provides the instantiation of the first, showing how natural selection ‘runs’ here on earth. Similarly, the proposal in the case of our seventeenth century gentlemen is that while members of any group must trust each other, the actual mechanism by which gentlemen managed this task is provided by the details that Shapin’s honor explanation provides.

I am on the whole sympathetic to this irenic solution to the problem of the relation between the two explanations, but there may still be some difficulty of epistemic competition. Purporting to explain the grounds of my love for another, I may cite her extraordinary attention to detail, her rigid moral standards, and the righteous indignation she feels at every slight, and then go on to say why these are such good traits. Knowing, however, that I am deeply narcissistic and that the traits I have cited are precisely those I share with my beloved, you may reasonably suspect that the real reason I love her is not that these traits are so desirable, but just that they are ones I share. Similarly, a sceptical reader of Shapin’s descriptions of the epistemic values invested by gentlemen in other gentlemen’s perceived freedom, virtue and honor may still leave with the impression that what is really managing the testimony in this case is simply the maxim, ‘trust your own kind’. Insofar as this is so, Shapin’s honour explanation may be right about the importance of the gentlemanly qualities he cites yet mislocate their import and exaggerate their moral aspect, since those qualities would function more as does a certain sort of accent than as genuine moral justifications.

Another criticism of Shapin’s particular explanation for gentlemanly trust moves in the opposite direction from my observations arising out of the communal explanation of trust. The ubiquity of testimony requires that we maintain especially strong trust relations within our own community, but it also requires that we trust outsiders. Shapin illustrates this in considerable detail, devoting a chapter to the trust that Boyle had to invest in the testimony of his technicians. Other ungentle folk upon whose testimony gentlemen depended, many also mentioned by Shapin, include women, servants, instrument-makers, apothecaries, printers, messengers, and on and on. But this obvious and pervasive dependence on outsiders’ testimony seems strikingly at odds with the general message of the honor explanation: gentlemen only trust gentlemen. Shapin must be aware of this tension, but he says surprisingly little directly to reduce it. What he does emphasise is invisibility of trust in the ungentle, so that for example a technician’s observation reported to a gentleman is treated as the gentleman’s own observation. Shapin’s discussion here is extensive and nuanced, but for present purposes the main point is the claim that ungentle testimony was only trusted when vouched for by a gentleman, in which case it was treated as in effect the testimony of the gentleman himself, and thus made trustworthy for other gentlemen (Shapin, 1994, pp. 400–401). The role of vouching in the epistemology of testimony is an interesting issue, and receives in this book an interesting treatment, but it does not remove the tension I have flagged. For as

Wittgenstein never said, all vouching comes to an end. Boyle could vouch for his servants' results, but he often had to accept them himself without any other gentleman's voucher, and of course all other scientists were in a parallel situation. Moreover, even the knowledge of who is a gentleman must have depended in part on ungentle testimony. So, however important the role of freedom, honor and gentlemanly virtue, these traits cannot fully explain how gentlemen managed their judgements of credibility.

I will end this section by flagging a different but fundamental limitation on the scope of the honour explanation, one which would apply with equal force to a communal explanation. These sorts of account make credibility remarkably independent of the content of the testimony. What counts is who says it, not what is said. In this sense, these accounts are indiscriminating: they do not discriminate among different assertions of the same actor. Yet of course all of us do make such discriminations. We trust a person on some matters but not on others, and on a particular matter will accept claims only within a certain range of plausibility. Shapin's lack of focus on this sort of discrimination may be related to his interest in honesty rather than competence. To some extent, judgements of honesty attach to the individual as a whole, though of course even here we think that someone is much more likely to lie on some topics than on others. Competence, however, is more obviously connected to the content of the assertion: a person is not judged competent in all matters, but only in this or that area. In any event, it is clear that the decision whether to believe someone depends not just on who they are but also on what they say and how what they say fits with what the audience already accepts. The central question about testimony is not just whom to trust, but what to believe.

5. The Humean Predicament

Like *A Social History of Truth*, Coady's *Testimony* does a wonderful job of bringing out the ubiquity of testimony and is particularly good on the hard case of observation. For Coady, the central point about observation is not just that it is almost always someone else's, but that even personal perception has no straightforward epistemic priority over testimony as a source of knowledge. One cannot enjoy the testimony of others without perceiving it but, as Coady argues at length, this entails no general epistemic priority. As I mentioned at the beginning of this essay, testimony may influence perception, insofar as perception is laden by theories themselves held on the word of others. More straightforwardly, we may in particular cases rightly put more weight on what other eyewitnesses tell us than on what we ourselves thought we saw. A perception is not always worth a thousand words.

The epistemic inter-penetration of testimony and perception provides Coady with a powerful weapon against reductive accounts of testimony in general, and David Hume's account in particular, as it appears in 'On Miracles', Chapter 10 of his *Enquiry Concerning Human Understanding* (Hume, 1748). The main aim of that

chapter is to show that testimony can never provide a good reason to believe in miracles. Two of Hume's arguments are particularly memorable. The first is that, since they are by definition violations of laws of nature, miracles are always epistemically undermined by the available evidence, since the laws of nature are just those generalisations most strongly supported by the course of our experience. The second is that it would never be rational to believe a miracle on the basis of testimony unless the falsity of the testimony would require a greater miracle.

Coady's fourth chapter, 'Testimony, Observation, and the Reductive Approach', is an extended attack on Hume's account of testimony, but it does not focus on these two famous arguments. Instead, it attacks the general epistemology of testimony that emerges from Hume's discussion, an epistemology that, for Coady, epitomises the wrongheaded reductive approach. To his credit, Hume is not one of those philosophers who denies the importance of testimony: 'there is no species of reasoning more common, more useful, and even necessary to human life, than that which is derived from the testimony of men...' (Hume, 1748, p. 74). That could have been written by Coady. But Hume also seems to hold that any warrant we have for believing a particular piece of testimony must rest on some sort of enumerative induction: 'The reason, why we place any credit in witnesses and historians, is not derived from any *connexion*, which we perceive *a priori*, between testimony and reality, but because we are accustomed to find a conformity between them.' (p. 75). It is this idea that our warrant for testimony reduces to the evidence of enumerative induction that Coady sets out to demolish.

Coady's attack is relentless, elaborate, and powerful. Hume's proposal seems to be that you ought only to believe a person's claim to have observed a particular situation if that claim-situation pair is an instance of a generalisation linking types of claims with types of situations, where you have independently observed many other instances. Coady's attack can be seen to consist of three strikes. The first is to show how Hume's discussion systematically hides the fact that our evidential base is far too slender to underwrite in this way even a small fraction of the testimony we rightly accept. Perhaps the main device Hume uses here is to appeal to the correlations *we* have observed to obtain between various types of testimony and the facts. This appeal to communal observation closes a vicious circle, since you can only in general know what others have observed on the basis of *their* testimony. The only evidence that you can legitimately appeal to consists of correlations between what you yourself have heard and what you yourself have seen, and this provides far less evidence than would be required to support inductively the wide range and variety of generalisations that would cover all the unchecked testimony you actually accept.

Coady's second strike hits the structure of the generalisations that Hume needs. These must link types of report with types of fact; but what are the relevant types? Coady distinguishes two ways of classifying report: by reporter (Shapin's preferred taxonomy), or by the content of the report. In either case, whether or not something

is a report of the relevant type must be something that can be determined by personal observation: it must not itself rest on testimony. This rules out at least some of the natural reporter types. Thus, you cannot look for connections between the facts and the testimony of experts in a particular area, because you cannot tell by personal observation whether someone is an expert. Indeed you cannot even retreat to the type of 'so-called expert', since even this is a communal matter. (Presumably Coady would make similar claims about the reporter type 'gentleman'.) Coady presses a different sort of objection to a taxonomy of reports by content. The difficulty here is that content can be classified in too many ways. One and the same report may, for example, be of the empirical, medical, geographical or existential type. The same report will thus instantiate many different testimony generalisations, some of which will be supported by personal observation, some disconfirmed and, harking back to the first strike, most simply untested. In other words, the Humean proposal runs the risk of generating inductive inconsistencies, blocking any inference whatever.

Coady's third strike, and the one on which he puts the greatest weight, attempts to show that Hume's reductive approach is fundamentally incoherent. Coady takes it that Hume's position entails the possibility that your personal observational research might reveal the awful fact that there is no correlation whatever between what people tell you and what is the case. In other words, you might discover that everything you are told is false. Coady uses the claim that this is possible as a *reductio ad absurdum* of Hume's approach. The idea that all reports are false is incoherent, according to Coady, for three reasons. The first is that in this environment the utterances in question simply would not be reports. They would have to be speech acts of some other type, if they were speech acts at all. Secondly, even if we supposed that the speech act of reporting could somehow survive in the environment of blanket falsehood, you could never check the correlation between reports and the facts, since you would have no way of determining what the reports *mean*. Finally, the supposition that all testimony is false would make it impossible for anyone to learn their language, and indeed the very possibility of a public language is in jeopardy.

The structure of Coady's third strike is reminiscent of Kant's claims about the impossibility of universal lying in the context of the categorical imperative, a comparison that has already served us in connection with Shapin's claims about the moral texture of testimony. Kant says, '... I could will the lie, but not a universal law to lie. For with such a law there would be no promises at all... Thus my maxim would necessarily destroy itself as soon as it was made a universal law.' (Kant, 1785, First Section). Similarly, Coady's claim is that while of course some testimony can be false, there is no possibility of the universal falsity of testimony, since in such a situation there could properly be no testimony at all.

I find Coady's general case against Hume exceptionally interesting and also disturbing, in part because Hume is my favourite philosopher. The dark jewel of

Hume's epistemology is his sceptical argument against induction and, while I cannot accept that argument, my largely fruitless efforts to refute it have provided a primary focus for my own philosophical work. Nevertheless, as many commentators have noted, Hume's extraordinary case against the justification of any form of non-demonstrative reasoning sits alongside a descriptive account of our actual inductive practices that is much less impressive. That account is Pavlovian, an appeal to 'custom and habit' that would reduce all non-demonstrative inferences to expectations that follow the principle 'More of the Same'. This account, it seems to me, does justice to almost no aspect of our inductive practice, so it is perhaps not surprising (if only itself on simplistic enumerative grounds) that the Pavlovian account also serves Hume ill when applied to inferences from testimony. This weakness of Hume's general account of our inductive practices may also however create a weakness in Coady's own line of argument, for it seems to me that Coady is too quick to move from the failure of Hume's attempt to reduce testimony to primitive enumerative induction to the conclusion that there is no general account of inductive warrant that might subsume inferences to the reliability of testimony. That is a point I will return to in the final section of this essay; now, however, I want to see whether there is anything at all that can be said to blunt Coady's specific attack on Hume.

I will take Coady's three strikes in order. As for the first, which accuses Hume of conflating personal and communal observation, I'm afraid Coady scores a palpable hit. Hume does often seem to construe observation communally, as when he writes, '... it is a miracle that a dead man should come to life; because that has never been observed in any age or country. There must, therefore, be a uniform experience against every miraculous event...' (Hume, 1748, p. 77). One might plead that this mistake illustrates the remarkable power of an effect that Shapin has discussed in detail: the invisibility of testimony. Once testimony is accepted, on whatever grounds or none, we very often treat the information as if it were a product of personal observation. This even a Hume seems to do, and even when he has the greatest philosophical motive not to. He behaves rather like the logician who wrote to Bertrand Russell, 'saying that she was a solipsist, and was surprised that there were no others' (Russell, 1948, p. 196).

One might attempt to defend, not just to excuse, Hume's appeal to communal observation, by taking him to have telescoped what is in fact a legitimate transition from the personal to the communal. The idea, in its crudest form, is this. You begin by establishing the veracity of your acquaintances, using only the resources of your personal observation. Having done this, however, what you credit are not just their reports of what they have themselves observed, but also their own epistemic standards for crediting the testimony of yet other people, most of whom you will never meet and most of whose reports concern states of affairs you can never directly verify. (A more refined version of this scheme would have you distinguish between an acquaintance's track record for reports of what she has observed and

her record for reports of what she accepts on testimony.) In effect, you come on the basis of personal observation to treat other people's observations and inferences as extensions of your own. They will act likewise for their acquaintances, and the result will be an expanding sphere of trust.

Coady can perhaps be criticised for apparently not considering any such telescope defence of Hume's appeal to communal observation, but I suspect that no such defence will succeed in the end, though it is not immediately obvious why not. I suppose the main reason for my scepticism is that I have been convinced by Coady's general case for the ubiquity of testimony and for the corresponding testimony-laden feature of almost all our beliefs. This leaves the Humean in the unhappy position of having to justify an enormous number of beliefs on an extremely slender evidential base. The telescope defence attempts to manage this by a kind of amplification of inference, so that the few direct correlations you can establish between the first-hand testimony of your acquaintances and what you have yourself observed are used to support the massive extrapolation you need in order to accept the testimony of so many strangers. It seems a precarious inference.

Coady's second strike—the objections from classification—seem to me to have more uneven impact. One of these objections is another application of the testimony-ladenness of belief: we cannot test correlations between experts, say, and the facts, because we cannot determine who is an expert without appeal to testimony. The general point also applies to the fact side of the generalisations: the facts in question will very often not be possible to determine without relying on testimony. I am somewhat less impressed by another of the objections from classification, the objection from multiple taxonomies. Coady's observation that the same testimony may be classified in many different ways is a point that applies to observation generally. This obviously does not make inductive inferences impossible, however difficult philosophers have found it to describe how the problem of inductive inconsistency is avoided in practice. (The so-called New Riddle of Induction, originally due to Nelson Goodman (1953, Chapter III), is an example of an especially general and recalcitrant version of this difficulty.) It may be that the problem of divergent classification has a special aspect in the case of testimony that makes it vicious here when it is not in other cases, but Coady does not attempt to show this.

I turn now to Coady's third strike against Hume. Hume's reductive approach is claimed to be self-defeating, since it must assume the possibility that there is no correlation between testimony and the facts (the possibility that all testimony is false), but this possibility would preclude the existence of a public language. This argument is extremely interesting but also problematic: the problems cluster around the status of the assumption that we might discover by empirical means that all testimony is false. There are two questions I want to consider. The first is whether the Humean project really does require the possibility of global falsity; the second is whether the idea of discovering global falsity by empirical means is incoherent.

Coady takes Hume's remark about the absence of a perceived *a priori* connection between testimony and reality as evidence that Hume accepts the possibility of global falsity. Coady also seems to suppose it to be clear that Hume's reductive project requires this assumption. As Coady puts it, '... on my construal of Hume's programme, his argument is defeated if to understand a language requires the admission of only one true report expressed in it.' (Coady, 1992, p. 152) But I am not convinced that Hume's programme does require the assumption. Let us suppose for the moment that the impossibility of global falsity would indeed be fatal to Hume's project. If this were so, Coady could have shortened his chapter considerably, since it is obvious that, given what testimony there is, global falsity is impossible, for at least two simple reasons. There seem to be certain testimonial acts which behave like performative utterances in a way that pretty much rules out falsehood. Examples might include 'No Smoking' and 'Meeting Point' signs, as well as the reflexive case of the 'I am speaking now' utterance. Secondly, the testimony you receive will include pairs of claims that are strict contradictories, for example when one person reports something and another claims that what was just reported was false. From this it follows that not everything you have been told could be false. If Hume really needs the possibility of global falsity, he is out of luck.

Hume is very unlikely to have been unaware of this point about contradictories, particularly as his discussion of miracles emphasises 'the opposition of contrary testimony' (p. 75). Nor would he in my view have been bothered by it, since the reductive project does not require the possibility of global falsity. The impossibility that everything you are told is false, because you are told contradictory things, in itself raises no barrier to the project of justifying any testimony you accept in terms that do not already assume the truth of any testimony. If this isn't obvious, perhaps the following analogy will help. Suppose that you are concerned with the very limited question of the veracity of two witnesses at a trial. Here you have no interest in the general justification of testimony, but you decide to assess the credibility of these witnesses in terms that do not already assume the truth of any of their own testimony, however much it may rest on the testimony of others. This miniature reductive project could well succeed even though the testimonial output of the two witnesses includes two claims, one of which is the negation of the other. Of course the fully general reduction of testimony that Hume envisages has much less chance of success than this highly restricted project, for the good reasons that Coady's book brings out—but the impossibility of global falsity is not among them.

Hume's project does not require the possibility of global falsity; but does it require something close, say that we might discover that most testimony is false? Even this weaker assumption might be strong enough for Coady's *reductio*, but it may still be stronger than the reductionist requires. Indeed it seems to me that the question of how reliable testimony must be is independent of the question of whether a reductive programme can succeed. Perhaps another analogy will serve

here. Suppose that I am an optimistic Cartesian, optimistic because I have convinced myself in some way that most of my perceptual beliefs must be correct, yet committed to a reductive justification of perceptual belief. I have thus convinced myself that demon-scepticism, where almost all my perceptual beliefs would be false, is incoherent. I may nevertheless attempt to provide a justification for each of those beliefs in a way that does not rest on perception. Moreover, I may have a motive for this project, since while I take it that they could not all be false, each belief, taken singly, could be. Similarly, I might have convinced myself that most of the items of testimony I receive must be true, yet still look for a reductionist justification for each of them. When Hume claimed that there is no *a priori* connection between testimony and reality, I think he just meant that in general 'Says P' does not entail 'P'. This is true, and it makes sense of the reductive project of justifying testimony in non-testimonial terms.

I turn now to my second question: has Coady shown that the Humean method of correlation could not yield the judgement of little or no correlation between testimony and the facts? Imagine that you begin to check those bits of testimony you can, and the bad news starts pouring in: in each case, what you have been told turns out to be false. Coady's idea is that, as you started to extrapolate to the general unreliability of testimony, you would simultaneously undermine your judgements about the content of what you had been told and indeed that what was said was testimony at all. At the limit, what you would have would not be a practice of unreliable testimony, but no assertion at all.

I will consider whether it really is impossible to have testimony that is generally false below, when I consider Coady's positive proposals for the vindication of testimony; for now, I want to make a different point. Even if Coady is right that, strictly speaking, the notion of a discovery that there is lots of testimony but it is all or mostly false is incoherent, he has not shown the incoherence of a very similar situation, where Humean testing leads you to judge that there is no reliable testimony. In this Kafkaesque situation, you come to the realisation that you cannot make sense of what people are telling you, just because the apparent falsity of everything they say undermines your judgement that they are in fact speaking a language you can understand, or indeed any language at all. It is quite unclear what the difference is, for present purposes, between the possibility that you might discover a complete lack of correlation between testimony and reality and that you might discover that there is no testimony. I will lean yet again on analogy. Consider the generalisation that electrons have negative charge. If this is taken to be non-definitional, we might find that the correlation breaks down on investigation. If, on the other hand, we take the claim to be a partial definition of what it is to be an electron, then perhaps it is incoherent to suppose that the correlation might fail: if it doesn't have negative charge, then it just is not an electron. But the same empirical investigations could yield a parallel judgement, namely that there are no electrons.

There is one more observation I would like to make about Coady's attack on Hume, before turning to his constructive use of the argument from language. If Coady's arguments were sound, then he would have established not only that a Humean reduction of testimony is impossible for us, but something much stronger, namely that there could not exist a Humean community, a community of language users for whom a Humean reduction of testimony was possible. For his argument aims to show that reductionism entails the possibility of global falsity, and the mere existence of testimony rules out this possibility. Coady has convinced me, primarily by means of his first strike, which emphasised the extreme narrowness of the empirical base, uninfected by testimony, that a Humean would have to employ for a reduction, that we are not a Humean community. But it still seems coherent to imagine a community who were much more cautious about testimony or much better at first-hand observation than we are, whose members could justify each bit of testimony they accept in terms that did not themselves rest on testimony. Perhaps in the end this is not really possible, but Coady's arguments do not succeed in showing this.

6. The Vindication of Charity

In Chapter 9, 'Language and Mind', Coady attempts a non-reductive vindication of the reliability of testimony. His general strategy is a boosted inversion of his third strike against Hume in Chapter 4. There he argued that the unreliability of testimony would entail the impossibility of language; here he tries to show how the existence of language might be used to provide an argument that testimony is generally reliable. Even if one accepted Coady's linguistic argument in Chapter 4, however, its simple inversion would produce only a very modest result. The earlier argument only aimed to show that the possibility that all testimony is false would entail that there is no language, so a simple inversion would only show that at least one item of testimony is true. So considerable bolstering of the argument is required if it is to yield the desired result—bolstering which Coady duly supplies.

The main additions are two. First, Coady brings in arguments due to Donald Davidson that a correct understanding of the semantics of language entails that most of a language user's beliefs must be true. Coady does not fully endorse Davidson's argument, but he does take it to show that we must treat most testimony as true. Second, Coady appeals to the way different sources of information, for example testimonial and perceptual, combine to generate a belief ('cohesion') and the way beliefs from different sources may support each other ('coherence') as a means of vindicating the reliability of the testimonial sources of belief. The aim of the exercise is not to show how testimony can be justified in testimony-free terms, but to show how our practice of accepting testimony can be seen to be non-circularly self-supporting.

Davidson's arguments are subtle and difficult but, with some trepidation, I will comment on the use Coady makes of them. Crudely put, Davidson argues that an

interpretation of a speaker that would make most of her beliefs false, by our lights, is *ipso facto* an incorrect interpretation. It follows that most of the beliefs we attribute to others must agree with our own. Davidson then makes the transition from agreement to truth by means of the device of an 'omniscient interpreter'. Such a being has only true beliefs, so when he finds himself in broad agreement with the beliefs of those he interprets, as he must, those beliefs will also be true (cf. Davidson, 1984, especially Essay 9, 'Radical Interpretation'). Coady rejects the step from agreement to truth, in part because of some unclarity he exposes in the very notion of an omniscient interpreter. (Coady asks, for example: if the omniscient interpreter is so smart, why does he need to engage in the process of interpretation in the first place?) I share Coady's scepticism about the argument from agreement to truth. Why not suppose instead a dim interpreter, most of whose beliefs are false? In this case, agreement between native and interpreter would entail broad error. So the device of the omniscient interpreter appears subtly to beg the question.

Coady does, however, more-or-less accept the first stage of Davidson's argument, the argument to the necessity of broadly shared beliefs, and he takes this to suggest that we could not in any case find that most testimony is false, and therefore that testimony must be reasonably reliable. I have several inchoate objections to Davidson's first argument, but I will here just mention two objections peculiar to the use Coady makes of it. The first has been raised by Elizabeth Fricker (1995, p. 410), in an excellent essay that has substantially aided my understanding of Coady's arguments and of the epistemology of testimony generally. According to Davidson, those of the agent's beliefs we must agree with are so obvious as not to be worth asserting: Davidson does not claim that we must agree with most of the agent's *assertions*. So it does not follow from Davidson's argument, at least as Davidson himself presents it, that we must agree with most of the agent's testimony. Secondly, since Coady does not accept Davidson's omniscient inference from agreement to truth, even if Coady's argument succeeded in showing that you had to agree with most of the testimony you received, it would not follow that testimony is generally reliable in the sense of generally true. Our epistemic situation in the case of testimony would have been shown to be no more secure than would be our situation in the case of perception, if all that one could show about perceptual beliefs is that they are generally involuntary. If we could not help but have these beliefs, we could not *find* that they were false, but they might well be false just the same.

I am much more sympathetic to the other strand of Coady's vindication of testimony, the argument from cohesion and coherence. Coady appeals in part to the work of Tyler Burge and Hilary Putnam on 'externalist' theories of meaning, according to which the content of language and thought is determined by factors outside the mind of the speaker or thinker. This general point of view fits well with Coady's general anti-individualist perspective on cognition, but its specific

bearing on the veracity of testimony is not developed in any detail. But Coady also appeals to ways we rightly gain confidence in the truth of an item of testimony if we find it to fit together in the right sort of way with other testimony we receive and the deliverances of perception, memory and inference. Coady is surely right to claim that these considerations of coherence provide a warrant for the testimony we accept, though he is also right to say that this gives no support to the Humean project of justifying the truth of particular reports by appeal exclusively to evidence that can be known without relying on the word of others.

One of the striking features of Coady's book, however, is how little it has to say about the warrant for testimony. In particular, and in striking contrast to Shapin's discussion, Coady rarely addresses the question of how we actually discriminate between the testimony we accept and the testimony we reject. For much though not all of Coady's discussion, it is as if trust is almost entirely blind: we simply accept whatever we are told. Given the ubiquity of testimony, perhaps Coady is right to suggest that most testimony is accepted as a matter of course, but even if conscious deliberation over whether or not to believe something one is told occurs only relatively infrequently, it still occurs all the time. Both the conscious and the unconscious discriminations we make raise some of the most interesting issues in the epistemology of testimony. In reading his book, it sometimes seemed to me that Coady downplayed these issues for fear of giving false comfort to the reductionists but, as Coady's own insights make clear, the implausibility of any Humean reduction of testimony is quite compatible with a general account of the epistemic discriminations we all make among the varied and often contradictory pieces of testimony we receive. In the final section of this essay, I want to suggest one direction one might look for such an account.

7. Testimony and the Best Explanation

There are two different sorts of account of the epistemology of testimony we might seek, both of which might be considered reductive. The first, which we might call a 'premise-reductive' account, is the project that Coady attacks, an attempt to show how every testimonial belief the agent is warranted to accept can be justified in terms that do not themselves appeal to beliefs based on testimony. Coady's attack on this project is effective, primarily because he shows that the base of premises available for such justifications is too small for the job. If one wished heroically to try to save this project, one might begin by rejecting the limited scheme of enumerative induction that Coady takes from Hume in favour of a broader and more realistic account of inductive inference. Still, in my view there is little reason to suppose the premise-reductive project to be either possible or particularly desirable. Our actual justificatory practices are certainly not of this form: in deciding whether to believe an assertion, we do not limit our consideration to things known independently of testimony, and the non-reductive justifications we do supply are not circular. Moreover, the rules of inference we actually use

are themselves partly constituted by background beliefs that would not be available in the extraordinarily austere environment of a testimony-free evidential base. If it is possible at all, the premise-reductive project would be a radical and highly artificial reconstruction of our knowledge of dubious epistemic value.

At the same time, all of us have the need to 'manage' testimony that Shapin illuminates. This need to determine which assertions of others are belief-worthy strongly suggests that we must deploy some principles of justification (which is not to say that these principles are the same for all societies). I have argued that these principles cannot be purely ethical and cannot simply pick out trustworthy individuals; the principles must be cognitive and sensitive to the content of the testimony on offer. And it may be that these principles of justification or inference may apply to testimony without being peculiar to it. That is, even if a premise-reduction of testimony is undesirable or impossible, we may be able to provide what we might call a 'rule-reductive' account, an account which would show that warranted testimonial beliefs are based on rules of inference or mechanisms of belief acquisition that apply to the beliefs from various sources, not just the source of testimony. As I read him, Coady does not deny the possibility of this sort of rule-reduction, but nor does he go very far towards suggesting what such an account might look like. And in spite of Shapin's emphasis on the problem of discrimination, his emphasis on trust seems unfriendly to a rule-reductive approach, since a rule which relies essentially on the moral status of the source of information appears only applicable to human sources, that is to testimony.

What are the prospects for a general account of the epistemology of testimony, in this rule-reductive sense? If all testimonial beliefs are inferential, then they are broadly inductive since, as Hume maintained, the fact that someone tells you that something is the case almost never entails that it is. A rule-reductive account would then consist in a general account of inductive reasoning, applicable to non-demonstrative inferences generally, though without any presumption that inferences from testimony could ever in general be reduced to inferences from testimony-free premises. Retaining for the moment the assumption that testimonial beliefs are inferential, there seems to be a strong presumption in favour of the existence of such an account, on the grounds of the role of cohesion and coherence in belief management. As Coady stresses, the testimony we accept is strongly supported by the way the uptake of testimony depends upon an integration of testimonial and other sources, and the way the beliefs thus acquired cohere with beliefs acquired by other cognitive routes. This holistic mechanism strongly suggests that, insofar as acquisition of belief from testimony is based on inference, the rules of inference are not specific to testimony.

A parallel case could be made for the other sources of belief, such as perception, memory, and further inferences drawn from these sources, to show that they too must rely on rules of inference that are not specific to those modalities. But the case for an account of testimonial inference that appeals to more general rules of

inference is particularly strong because, unlike other sources, testimony is a virtually universal source: almost anything you can believe at all and almost anything you can believe on the basis of another source, you can also believe on the basis of testimony. This makes the integration of testimonial and other beliefs particularly extensive.

But are testimonial beliefs inferential? Coady often writes as if they are not: we are surrounded by people saying various things, and we just thereby come to have the corresponding beliefs, without deliberation. The general rule-reductive project need not, however, depend on the claim that testimonial belief is all or mostly inferential. Whether inferential or not, we have mechanisms for acquiring our beliefs, and the rule-reductive idea is simply that these mechanisms will be at least largely inter-modal. Indeed the very distinction between direct and inferential belief is not all that clear. Moreover, insofar as we can make a relevant distinction here, testimonial beliefs seem to me rather more inferential than Coady portrays them to be. One way of seeing this is to note that the integration between beliefs from different sources is not simply a discovery, it is something we create. Coherence is a constraint on belief acquisition, pushing us to be selective about which testimony we accept, not just a happy outcome of blind trust, and this gives the testimonial mechanisms at least some of the characteristics of inference.

Fricker (1995, pp. 404–406) has strengthened this point, by emphasising the extent to which our apparent gullibility is actually underpinned by a steady mechanism of often sub-conscious monitoring for plausibility. When I ask the ticket-seller at the rail station about the price of a return ticket to Sheffield and she tells me it is thirty-six pounds, it may appear that I simply and directly believe whatever I am told, but this is not the case, since I would not have believed her had she said thirty-six pence or thirty-six thousand pounds. Testimonial beliefs are not in general like scientific hypotheses or even like the detective's inferences based on the testimony of a dubious witness, but they may yet have a strong inferential component. One difference between mundane cases of testimony and some of the more obviously inferential activities of the scientist or the detective is that scientists and detectives sometimes have to do considerable work to come up with the explanations they will go on to infer, whereas in mundane testimony the belief we come to acquire is given to us on a plate, since it is simply the content of the testimony itself. Nevertheless, our selectivity, actual and counterfactual, about what testimony we will accept and what we will not shows that this real difference does not make testimonial beliefs non-inferential.

The fact that I would not believe that the rail ticket costs thousands of pounds does not show that my acceptance of the more modest testimony of the ticket-seller was based on a deliberate inference. After all, I would not believe my eyes if I started to see large pink elephants dancing around my office either, but this does not show that my actual perception is based on deliberate inference. Nevertheless, Fricker's point about counterfactual monitoring suggests an attractive 'default-trig-

ger' picture of testimony, where the default situation is often simply to accept what we are told, but this may be overridden by a trigger that switches us into inferential mode. These triggers will take various forms. Some will concern types of people or situations that encourage healthy suspicion. Others will lean on the content, such as contradictions or lesser tensions between the testimony in question and the listener's background beliefs. Content may trigger an inferential mode in other ways as well; for example claims on controversial topics are taken not to be the sort that one simply takes on trust. There is much more to be said on this subject, but to object to a rule-reductive account on the grounds that testimony is non-inferential seems misplaced.

Another and related source of what is in my view a misguided scepticism about the rule-reductive project stems from an over-narrow conception of inductive inference that I have already flagged, which equates induction with the sort of enumerative correlation that Hume emphasised and that, for the case of testimonial inference, Coady demolishes. The proposal that 'horizontal' enumerative induction might be the general principle of inference that underlies diverse sources of belief is hopeless. This is obvious in cases such as perception, where it is clear that we can never be in a position independently to observe both our perceptual experiences and the external world in order to establish a correlation. Testimony is a much more seductive case, since it seems that we can in some cases independently check what we have been told without relying on yet further testimony, and this may fuel our hope, as it did Hume's, that a general justification of testimonial beliefs along these lines would be possible. The case is similar though much more extreme in the case of the problem of other minds. I can observe a correlation between behaviour and mental states in my own case, so I might be tempted, as some other philosophers have been, to suppose that I can provide an enumerative justification for my inference from other people's behaviour to their mental states. This inference, however, is obviously hopeless, rather like inferring that since the only oyster I could ever open has a pearl, all the others have pearls inside as well. The testimony case is less obviously hopeless but, as Coady shows, it is hopeless nevertheless. It does not at all follow from this, however, that no inductive justification is possible. The kind of induction we need is of the 'vertical' variety, where we infer from what we observe to what we have not observed and often could not observe. This is of course not peculiar to testimony, since it applies to any justification of beliefs in other minds, as it does to scientific inferences from data to hypotheses that traffic in unobservable entities and processes.

It is one thing to show the plausibility of the existence of some general mechanisms of induction that would apply to epistemology of testimony, quite another even to sketch a plausible account of what those mechanisms are. The main lesson of the last one hundred and fifty years of work on this descriptive project is that it is far harder than anyone would have expected. Even toylike accounts, whose range of application is restricted to the simplest cases, fall to pieces under investi-

gation. Thus innocent-looking accounts of the bearing of scientific evidence on theory, such as various versions of hypothetico-deductivism, are shown to have the absurd result that every observation supports every hypothesis. Given the absence of any even roughly adequate account of induction in other areas, it would be very difficult for anyone to show, against what I have been claiming, that testimonial inferences are *sui generis*. That however is pretty cold comfort, and it seems to me that the descriptive project remains well worth pursuing, and in a way that would incorporate testimonial inferences.

The account we want will apply to both testimonial and non-testimonial sources, and to their integration: that is what will make it rule-reductive. It must be 'vertical' and will need to be able to take on board the great heterogeneity of factors that go into the evaluation of testimony. For example, we want an account that gives proper weight both to those considerations that bear on the honesty of our informants and to those that bear on their competence. One familiar sort of account that seems to me promising in these respects is 'Inference to the Best Explanation', according to which we infer what would, if correct, provide the best explanation of our evidence (Lipton, 1991). This account faces various difficulties (I prefer the word 'challenges'), but I want to end this essay by suggesting some of the reasons why it may nevertheless be well suited to the integrative project of bringing testimonial inference into the fold of a general account of belief acquisition.

Inference to the Best Explanation makes plenty of room for vertical inference. When scientists infer a hypothesis that traffics in unobservable entities and processes, it is hopeless to characterise this in horizontal, enumerative terms, but natural to say that the hypothesis was inferred because it would provide the best explanation of the available data. Similarly, insofar as it is legitimate to construe beliefs about the other minds and about the external world as inferences from behaviour and experiences respectively (which may not be very far), one may say that beliefs about mental states and medium-sized dry goods are held because they are taken to provide the best explanations of the relevant evidence, though here too enumerative induction is out of the question. In the case of testimony, the most straightforward application of Inference to the Best Explanation would be to say that the agent infers that what the informant said is true just in case the truth of what was said is (part of) the best explanation of (among other things) the fact that the informant said it (cf. Fricker, 1995, pp. 404–405).

Inference to the Best Explanation is well suited to take into account the role of coherence considerations in the evaluation of testimony. Your decision on whether or not to accept a piece of testimony depends on how well it would fit together with the other beliefs you hold, and the notion of fit can be articulated in explanatory terms. Roughly speaking, the question is whether the candidate belief would, in the context of those other beliefs, provide the best explanation of, among other things, the informant's utterance. This aspect of Inference to the Best Explanation also suits it to the sort of account of inference that integrates testimonial and other

forms of belief acquisition. Coady's own comments on cohesion and coherence are naturally taken in this way, though his focus, as we have seen, is on the inference from cohesion and coherence of the corpus of beliefs already held to their truth, rather than on the decision of whether or not to add what the witness says to that corpus.

There is I think only one place in his book where Coady explicitly discusses Inference to the Best Explanation, at the end of his chapter on Russell (p. 119). His discussion there, prompted by comments of David Lewis, is brief and negative, raising two specific objections to the application of the account to testimonial inference. One is that the data that are to be explained, in the case of testimonial inference, are themselves dependent on testimony. This is correct: although the bare observation that the witness said a particular thing may not be interestingly testimony-laden, clearly many other of the hearer's beliefs that play a role in the explanatory evaluation are so laden. As Coady recognises, however, this is only an objection to the use of Inference to the Best Explanation in the context of what I have called a premise-reductive account, which I have joined Coady in abjuring. It is no objection to the rule-reductive project. Coady's other objection is that the criteria for 'best' in 'best explanation' are often unclear in such a basic epistemological context. If what Coady has in mind here is that it is very difficult for the philosopher to articulate the factors that lead people to judge one explanation better than another, then he understates the case: the criteria are unclear in virtually all contexts, not just in the basic epistemological ones. But this is in my view no reason not to work to articulate those criteria and, in any event, this problem of unclarity affects Coady's own appeal to coherence equally, since the coherence in question is at least often explanatory coherence which, like explanatory quality generally, is a matter of degree.

I have so far flagged two features of Inference to the Best Explanation that make it an attractive approach to the project of providing a rule-reductive account of the acquisition of testimonial beliefs: the first is that it sanctions vertical inferences, the second that it emphasises the role of coherence considerations in inference. These features correspond to two central aspects of testimonial inferences that are shared by many inductive inferences generally. But there are also certain features peculiar to testimony that Inference to the Best Explanation appears well placed to capture. One concerns the two quite different ways that an informant can fail us, mendacity and honest incompetence. Neither of these appear to correspond to the sources of error in other forms of inference, and the two are, at one level of analysis, very different from each other. Nevertheless, both fit naturally into the framework that Inference to the Best Explanation provides.

Our social intercourse is mediated by a continuous stream of inferences to explanations for the behaviour, especially the verbal behaviour, of those around us, and judgements about sincerity and competence are central features of the psychological accounts we generate to explain why people tell us what they do, accounts that

also draw freely on our beliefs about the non-social world and the place of our informants in it. When we explain why a person says something, the explanation need not rest on a prior determination whether what was said is true. Thus we may judge that the best explanation of her saying what she did is that she is unlikely to be deceived or deceiving on that sort of matter, without first independently knowing whether what she says is true. Consider a different sort of case, where we discover that someone has misinformed us. We often go on to attempt to explain this to ourselves, in terms of character, context and content, in terms of facts about people and facts about the rest of the world. The suggestion that we apply Inference to the Best Explanation to the epistemology of testimony is the suggestion that we often ask the same sorts of explanatory questions in advance too, in order to decide whether to accept a bit of testimony in the first place.

A further advantage of Inference to the Best Explanation is that it can account for the apparent tensions between different familiar maxims for the assessment of testimony. Shapin brings out the tension in an interesting discussion of typical Lockean lists of such maxims (Shapin, 1994, pp. 212–238). Shapin's strategy, reminiscent of Paul Feyerabend's (1975) discussion of counterinduction in *Against Method* (Chapter 2), is to find, for each testimonial maxim, a 'countermaxim'. Thus consistency between different informants is good, but may also be a suspicious sign of collusion or of a common and unreliable source; knowledgeable witnesses are good, but they can also be bad because they tend to over-interpret; confidently delivered testimony is good, but it is sometimes hesitant, stumbling testimony that should inspire the greatest confidence. Shapin draws two conclusions from the countermaxims he cites. One is that the principles for evaluating testimony must be taken to be uncoded skills and prudential maxims rather than formal epistemic rules. The other is that since, Shapin claims, 'the integrity of sources was the one inducement to assent which did not generate a countermaxim' (p. 238), a community's ability to locate honest, disinterested people was the fundamental means of ending debate about the truth of testimony.

Shapin's countermaxims are plausible and raise interesting issues, but it is unclear whether his conclusions follow. His second conclusion—his attempt to reduce the other maxims to the maxim of integrity, is part of his general project of reducing epistemology to ethics or social theory, an attempt I criticised above. The particular move here seems a clear example of Shapin's general suppression of incompetence and other forms of simple error. He mentions the possibility of honest mistakes, but then virtually ignores it. The maxim to trust honest people may have no countermaxim, but neither it seems does the maxim to trust people who know what they are talking about. Judgements about the honesty of the speakers do not always settle the question of the truth of what they say.

Shapin's first conclusion from the existence of countermaxims makes a better point. The standard maxims are not simple formal rules that allow for algorithmic application. As Shapin says (p. 239), they do not indicate how they are to be

balanced against each other, when a maxim is to be used and when its countermaxim, and their application depends crucially on background beliefs not specified in the maxims themselves. But maxims can, I think, very naturally be seen as rules of thumb whose warrant derives from the broader considerations that an Inference to the Best Explanation account would bring out. Whether we employ a maxim or its counter in a particular case depends on the overall explanatory account we favour for the facts about what was said, who said it, and how it was said, an account that will, as Shapin and I both emphasise, depend crucially on background belief. The factors that lead us to judge one explanation better than another in particular cases are themselves unlikely to lend themselves to formulation into neat maxims or be carried in the form of explicit knowledge in we who constantly deploy them, but this is no barrier to the intellectual project of elucidating their form. It is also worth noting that the general Inference to the Best Explanation scheme is compatible with (though it does not entail) the claim that Shapin and others have made, that the factors that induce credibility may vary from culture to culture, insofar as judgements of explanatory quality may so vary.

Shapin claims (p. 237) that his maxim to trust those with integrity and disinterestedness seems a tautology. Taken strictly, as a maxim concerning honesty but not competence, it is not tautologous at all, as the pseudomaxim 'trust the honest fool' shows. It is, however, easy enough to conflate the maxim of integrity with the genuinely tautologous statement that one ought to trust the trustworthy. A natural worry about Inference to the Best Explanation is that it has the same feature. Certainly if 'best' just means 'likeliest', the account veers towards the trivial, though even this excessively modest version may have some illuminating content, insofar as it brings out the links between testimonial and non-testimonial inferences, the need to construct a psychological interpretation of the speaker, and the essential role of background belief. Indeed, unlike the case of inference to scientific hypotheses, where the explanatory connection between what is inferred and the evidence is generally obvious, in the case of testimony it is perhaps not trivial even just to observe that what we believe on the basis of testimony is an explanation of why it was given. Nevertheless, a satisfying version of Inference to the Best Explanation will go further, cashing out 'best' in terms of the symptoms that guide our judgements of likelihood. Ideally, 'best' would be replaced by factors that have direct explanatory import, so that the account shows how we infer that the features of an explanation that would, if it were correct, make it the explanation that would provide the greatest understanding—the 'loveliest' explanation—are those that lead us to judge it also to be the explanation likeliest to be correct.

I have elsewhere attempted to make some progress on this general project in the case of non-testimonial inferences but, like pretty much everything to do with induction, the matter turns out to be very difficult. The project of articulating Inference to the Best Explanation for the case of testimonial inference may be even harder, but it still seems to me a lovely project, one that invites interdisciplinary

contributions from historians, philosophers, sociologists, cognitive scientists and others, and one in which Coady's and Shapin's work can play an important role. Their books make a signal contribution to our understanding of testimony. But don't take my word for it: read them for yourself.

Acknowledgements—I am grateful for very helpful discussions following talks I gave on the epistemology of testimony at the University of Cambridge, King's College London, Queen's University of Belfast, the University of Sheffield and Trinity College Dublin. For their illuminating comments on these occasions or in private I would like particularly to thank Arif Ahmed, David Bell, Jeremy Butterfield, Peter Carruthers, Marina Frasca-Spada, Donald Gillies, Chris Hookway, Nick Jardine, David Papineau, Simon Schaffer, Anne Secord, Jim Secord, Paul Teller, Ruth Weintraub and Kathleen Whalen.

References

- Coady, C. A. J. (1992) *Testimony* (Oxford: Oxford University Press).
- Davidson, D. (1984) *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press).
- Feingold, M. (1996) 'When Facts Matter', *Isis* **87**, 131–139.
- Feyerabend, P. (1975) *Against Method* (London: NLB).
- Fricker, E. (1995) 'Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony', *Mind* **104**, 393–411.
- Goodman, N. (1953) *Fact, Fiction, and Forecast*, (1973) 3rd edition (Indianapolis: Bobbs-Merrill).
- Hume, D. (1748) *An Enquiry Concerning Human Understanding*, E. Steinberg (ed.) (1977) (Indianapolis: Hackett Publishing Company).
- Kant, I. (1785) *Foundations of the Metaphysics of Morals*, L. W. Beck (trans.) (1959) (Indianapolis: Bobbs-Merrill).
- Lipton, P. (1991) *Inference to the Best Explanation* (London: Routledge).
- Locke, J. (1700) *An Essay Concerning Human Understanding*, 4th edition, P. H. Nidditch (ed.) (1975) (Oxford: Oxford University Press).
- Plato, *Theaetetus*.
- Russell, B. (1948) *Human Knowledge: Its Scope and Limits* (London: George Allen & Unwin).
- Shapin, S. (1994) *A Social History of Truth* (Chicago: University of Chicago Press).