

Knowing Through Being Told [Draft]

Alan Millar

1. *Introduction*

There are two extremes in the epistemology of testimony. One is represented by the following statement by Michael Welbourne.

All that is required of a listener who understands a knowledgeable teller if the knowledge is to be successfully transmitted to him is that he *believe* the teller.

(Welbourne 1986: 5.)

If this is right then if I know that you left your umbrella in my office, and I tell you that you have left your umbrella in my office, then you will know that you have left your umbrella in my office, provided only that you believe me and thus accept what I tell you. The other extreme is represented by a recent discussion of testimony by Richard Fumerton. Fumerton thinks it plausible that if I accept that P on Jones's saying that p then I rely, quite properly, on reasoning that can be represented as follows.

- 1) Jones said that p in conditions C (where C includes a description of Jones, his qualifications as an authority, and the circumstances under which he made the assertion).
- 2) People who assert that p in conditions C are usually saying something true.

Therefore,

- 3) p

(Adapted from Fumerton 2006: 80.)

On Welbourne's view acquiring knowledge from testimony can be easy. By contrast, Fumerton thinks that on the kind of model of justified belief that he favours 'the epistemic tasks one must complete in order to justify belief based on testimony are intimidating to say the least' (Fumerton 2006: 78).

Welbourne is surely right to suggest that we can readily acquire knowledge from testimony. You could easily come to know that you left your umbrella in my office on the basis of my telling you this. It is gross exaggeration to suggest that you would have to engage in intimidating tasks to acquire even a justified belief by this means. That said, Welbourne's statement looks like gross oversimplification. It overlooks the fact that recipients of testimony are discriminating and it seems obvious that they need to be so to acquire knowledge from testimony. It is likely to be relevant to your believing me, and to your acquiring the knowledge I transmit to you in the case of the lost umbrella, that you are aware of having recently having been in my office, that you know enough about me to know that I would be speaking from knowledge and would not be out to mislead you on such a matter. The view that Fumerton favours can be viewed as an attempt to take account of the various factors that are relevant to acquiring knowledge, or at least justified belief, from testimony. On this approach, justified belief in what one is told requires one to be able to identify those factors. It also requires one to be justified in accepting a generalisation connecting them to the truth of what is said (corresponding to premise 2 of his schema).

There is a position intermediate between the two just described according to which the appropriate default stance to testimony is belief. The view entails that one is justified in believing what one is told in the absence of reason not to.¹ It has the merit of accommodating the idea that recipients of testimony ought to be discriminating, at least to the extent of taking due account of whether there is reason to doubt what they are being told. But it is unclear just what one is supposed to do to meet this condition. Is it enough that nothing should strike one as providing reason not to believe what one is being told? Such a view would seem to sanction carelessness and gullibility. Is it necessary to establish that there is no reason to doubt what one is being told? Since the best way of doing that will be to have positive grounds trusting the speaker, it is unclear why we should introduce the notion of a default stance at all. Why don't we just focus on the epistemology of treating people as trustworthy?

There is no doubt that recipients of knowledge from testimony are discriminating; they don't believe just anybody on anything. Not should they. But I do not think that Fumerton's schema will do the job of explaining how we routinely reach the conclusion that what someone is telling us is true or of explaining why we can be right to do so. It distorts the phenomenology to suppose that we actually rely on inferences conforming to the schema. And to suppose that we would need to justify belief in a connecting generalisation in order to have a justified belief in what someone tells us makes gaining knowledge from testimony more elusive than it is.² I seek a way of accounting for how we are able to be sufficiently discriminating in our responses to testimony to acquire knowledge thereby, without commitment to the supposition that acquiring knowledge from testimony requires undertaking intimidating tasks.

The following gives an indication of some key features of the approach that I adopt.

- (a) The account applies to straightforward cases of the transmission of knowledge by telling. Here is an instance. Wondering if I have left a certain book in my study I phone home from the University. My son answers, looks in my study and sees that the book is on my desk. He tells me this. From his telling me I know that the book is on my desk. I take in that he is telling me that the book is there, and I accept that the book is there for the reason that I have been told this by him. So there is a minimal sense in which my knowledge is evidential—the evidence being my son's telling me that the book is on the desk. But this leaves out of account much of the relevant activity. It is possible that my son should tell me such a thing and that it be false. Sometimes he fools around, as when, to show up my gullibility, he tells me, falsely, that there is a smudge on my nose. It is not the bare fact of his telling me that the book is on my desk that does the trick. His telling me has features, including relational features, from which I tell that my book is on my study desk at home. Among those features are his manner of speaking, the content of the telling, the context—my having phoned home from the University and enquired about the book—and shared background knowledge, for instance, that if I am enquiring from the University for a book I

¹ For views of this kind, see Coady 1992, Burge, 1993, Welbourne 2001, Sosa 2006. Sometimes the view is expressed in terms of the idea that belief in testimony is *prima facie* justified. I take it that this amounts to the same thing. For trenchant criticism, see E. Fricker 1994 and 1995.

need I am unlikely to be in the mood for fooling around. In virtue of having these features and other features the telling can indicate that it is truthful. ‘Indicate’ has a strong sense. The telling can indicate that it is truthful in the sense that it would not be the telling that it is—a telling that has the various features that it has—were it not that it is truthful.

- (b) Much of our knowledge is perceptual, being gained through the exercise of perceptual-recognitional abilities.³ We may have, for instance, the ability to recognise that a sound is that of a fire alarm, or recognise by looking that the flowers in a vase are daffodils, or recognise that a fabric is silk from the way it feels. Possession of any such ability requires more than that the appropriate sensory system is functioning adequately; it requires that we have appropriate concepts *and* have learnt to apply them in response to what we perceive. We have to learn to recognise robins or goldfinches from the way they look. We have to learn to recognise people we know from the way they look or sound. Thanks to perceptual-recognitional abilities, perceptual knowledge can extend far beyond the superficial features of things that go to make up their appearance. We can recognise Bill and not just someone having the look of Bill. Such abilities involve refined sensibilities by which one is attuned to the significance of a host of physical cues that one would be hard put to describe. This is rather obviously so in the case of face recognition. The information we can capture by a description of a face is a tiny fraction of what we respond to when we recognise the person whose face it is. Perceptual-recognitional abilities clearly play a role in the epistemology of testimony. They figure at two different levels. The first level is concerned with how we recognise that we are being *told* this or that. The second level is concerned with how we tell from features of an act of telling that the telling is truthful. We come by knowledge that someone is telling us that *p*, not by inference from prior assumptions pertaining to the sounds the person uttered or the inscriptions the person has written, but simply through exercising an ability to recognise straight-off that the person is telling us that *p*. Focusing on recognitional abilities, here as elsewhere, opens up the prospect of explaining how a

² Fumerton is not troubled by the phenomenological problem (2006: 85), believing that the assumptions of the crucial inference can play an inferential role even if no conscious reasoning takes place. Even if that is right a problem for the view is whether the requisite generalisation is available. I discuss this below.

³ I put the notion of recognitional ability to work in Millar 2007 and 2008, which address issues arising from disjunctivism.

host of factors can be relevant to acquiring knowledge that something is so even though the knowledge does not result from reasoning from evidence pertaining to those factors. When you recognise that I am telling you that you left your umbrella in my office you are responding to features of my utterance and of the context that you would be hard put to articulate. But you do not infer from assumptions pertaining to those features that I am telling you that you left your umbrella in my room. You simply hear me to be telling you that. Our concern here, of course, is not just with knowledge that one is being told this or that, important though that is for the epistemology of testimony. It is with knowing that p through being told that p. The position taken in this essay is that such knowledge, in the straightforward cases of interest, is more like perceptual-recognitional knowledge than might at first appear. As I observed above, though there is a sense in which it is evidence-based our means of specifying the evidence does not suffice to articulate all that we respond to when we tell that p from being told that p. In this respect telling that someone's telling you that p is truthful is analogous to telling that deer have recently passed along a path from tracks on the path, or telling that it has rained recently from the wetness of streets and pavements. You need to be able to identify the right kind of tracks—the deer-indicating tracks—to tell from the tracks that deer have recently passed. You need to be able to identify the right kind of wetness—rain-indicating wetness—to tell from the wetness that it has recently rained. Similarly, I contend, you need to be able to identify an act of telling that p as a truthful telling to tell from that act that p. As already suggested, we sometimes tell from features of an utterance, not just that it is a telling that p, but that it is a truthful telling. In all of these analogous cases we tell from a situation that this or that is so. We perceive the situation, for instance, that in which there are tracks on the path—and we recognise the significance of the situation—that deer have recently passed. Talk of recognition in such cases is entirely appropriate. I contend that it is appropriate too in connection with knowledge of the truthfulness of an act of telling in the straightforward cases.

- (c) I place considerable emphasis on the idea that telling someone that something is so is a distinctive communicative act.⁴ My telling you that p is an act of saying to you that p by which I give you to understand that I am informing you that p. In so doing I give you to understand that I know that p (that I am speaking from knowledge) and intend that you should, by my saying to you that p, come to know that p. (For convenience, acts of saying that p should be taken to include cases in which the informant simply says, ‘Yes’, when asked whether p. Likewise, saying, ‘No’ in response to an enquiry whether p, should be taken to be an act of saying that not-p.) I shall invoke the idea that telling is a move in a *practice*—an essentially rule-governed activity or cluster of inter-related activities. The practice is that of transmitting knowledge. It is through initiation into the practice that we learn how to tell things, and how to recognise that we are being told this or that. But I shall argue that participating in the practice does not suffice to provide us with the means to recognise truthful acts of telling. For that we need knowledge of the informant that we do not have just in virtue of participating in, and understanding, the practice.
- (d) Much traditional epistemology obscures the importance of the distinction between, on the one hand, knowledge and, on the other hand, true belief that is justified but where the justification falls short of what is required for knowledge. This is especially evident in relation to knowledge based on evidence—that is, when it is knowledge that one thing is so on the basis of evidence supplied by something else. Evidential support admits of degrees. It is obscure just what degree of support is supposed to be required to provide grounds that are adequate for knowledge. Yet in practice it does matter to us whether we know or merely have grounds for some degree of confidence. For instance, by checking cinema listing I may come to know that a film I want to see is showing this evening at a local cinema. That is different from having a fairly high degree of confidence that the film is showing on the grounds that someone I know is pretty sure that he recalls seeing that it was listed as showing. The difference isn’t just that you can be confident in believing something false. Suppose that the film *is* showing this evening. If I have only a fairly high degree of confidence that this is so I am liable to take account of the possibility that it isn’t so. I may

⁴ For such a view, see Welbourne 1986 and 2001, and Moran 2006. I do not claim that the English term ‘telling’ and its cognates, when used in connection with declarative utterances, is used only for the kind of communicative act

try to determine for sure by looking for what the listings say on the matter. If that fails I am prepared to be disappointed and I would not vouch for the film's showing this evening. All this is independent of whether or not the belief is true. The distinction that matters for present purposes is between having grounds that settle the matter, so that in view of having those grounds one can be assured, and having evidence that merely warrants some fair degree of confidence. An account of the epistemology of testimony should explain how it can be that testimony that something is so can settle that it is so. My telling you that you have left your umbrella in my office clearly could settle the matter for you. Your reflecting on where you were when you last recall having your umbrella in your possession might well not. (Note that Fumerton's schema will not explain how the matter is settled since premise 2 does not serve to rule out the possibility that what Jones says is false.) Thinking of knowledge from testimony, in the cases of interest here, as involving recognition of the significance features of an act of telling, will help to satisfy the desideratum to account for how testimony that *p* can settle that *p*.

2. *Saying and telling*

Not every case in which people say things to us are ones in which they are telling us what they say. Telling someone that *p* is, as I just remarked, a distinctive communicative act. Central cases of the transmission of knowledge by telling are ones in which the recipients can easily pick up that the speaker is at least purporting to inform them that *p*. They deal in facts that the informant is clearly equipped to know. They do not deal in matters that cannot readily be settled. There are cases in which speakers 'sound off' about matters that are acknowledged to be matters of opinion or personal conviction, rather than the subject matter of information. People sounding off about Blair or Bush or climate change may mean what they say, yet not give it to be understood that they are informing those to whom they are speaking. The content of their utterances, as well as the context in which they are made, may mark them as expressions of opinion or conviction. A further interesting case of saying that is not telling is saying something by way of offering advice. Suppose that someone I know has had a breach with his parents. On learning of the circumstances I might suggest that he should try to put matters right with them by visiting them

that I am describing here. I shall reserve it for this purpose.

and talking about the problem. In saying what I do, I need not intend that he should take me to be *telling* him that he should put matters right. I am not informing him nor even giving him to understand that I am informing him. For one thing, the situation may be such that it is not easy to see what he should do. For another, even if I feel assured that my advice is sound I need have no intention to the effect that he should come to know what I say from my saying it. On the contrary, I may offer the advice as something he should think about and act upon if after thinking about it, he himself comes to view what I suggest as the right thing to do.

The question arises as to how we identify tellings and thus how we discriminate tellings from sayings that are not tellings. There is a perspective from which this can seem a formidable task. Against the background of the traditional problem of other minds, one might suppose that when people tell us something, and we understand them, we form a judgement that they are telling us such-and-such on the basis of evidence that can be specified independently of the meaning of the words used, and of their being used to tell us something. On this way of thinking, it is assumed that the evidence can comprise only what can be taken in perceptually, and that what can be taken in perceptually is confined to the level of superficial features. It does not include psychologically loaded facts to the effect that someone's words mean this or that or that he or she is telling us this or that.⁵ The main problem for such a view is that there is no reason to think that we routinely treat psychology-free facts as evidence for a judgement as to what the speaker is telling us. And if that is so, then such facts can hardly figure routinely in the explanation of how we can know things on their basis. A more plausible approach is to acknowledge that we can perceptually take in psychologically loaded facts.⁶ In the cases under consideration, we hear that the informant is telling us that such-and-such. We don't need to base our judgement on evidence any more than we do when we recognise a friend from his appearance.⁷

⁵ Fumerton thinks that the following is 'probably right': 'The meanings of sounds or marks are to be found in the head, or better in the minds, of those who produce those sounds and marks. To interpret reasonably those sounds or marks we need to solve the problem of other minds. ... When the language we are interpreting is our own—without begging the question, when the sounds or marks appear to be of the same type as we ourselves use to express thought—one might suppose that our best hope is to rely on some version of an argument from analogy or an argument to the best explanation' (2006: 83).

⁶ On this see McDowell 1981 and 1982. For related discussion of the scope of perceptual knowledge, see Millar 2000.

⁷ Here I respect J. L. Austin's (1962: 115) contrast between seeing the pig as opposed to finding evidence of its nearby presence.

How is it that we are able to take in ‘rich’ facts, and not merely facts concerning the superficial features of things? The answer is that we have appropriate perceptual-recognitional abilities. As I remarked in the previous section, such abilities enable us to have perceptual knowledge that extends well beyond the features that make up the superficial appearance of things to this or that sense modality. The exercise of such abilities is the acquisition of knowledge. The ability to recognise a goldfinch from its visual appearance—from the way it looks—is an ability to tell of something that it is a goldfinch from the way it looks. The ability is exercised only if I tell and thus come to know of the thing looked at that it is a goldfinch. For present purposes it is especially important to appreciate that the knowledge acquired in virtue of exercising such an ability is not implicitly evidence-based but is genuinely recognitional. Though I tell that the bird at which I am looking is a goldfinch from its appearance and may pick out some salient markings, it is the entire Gestalt presented by the bird to which I respond and which triggers recognition. Goldfinches have that distinctive flash of red on the side of the beak, but it is (showing you the illustration in the bird book) *that* distinctive flash of red. You need to see it to know what I am talking about. As I look at a goldfinch I do not propositionally represent its visible features in any very specific way, nor therefore do I work with a generalisation that connects specific features to being a goldfinch. No plausible model of perceptual knowledge of the presence of goldfinches should require this of me. Likewise, no model that represented my knowing that the shop assistant is telling me that if I buy two I get one free in terms of the application of a generalisation to the sounds that make up the utterance could do justice to the character of my ability to come to know such a thing. There is a host of cues by which I come to know such things. There are the sounds by which I grasp what is being said; I recognise the informant to be a shop assistant; the shop assistant speaks as if informing me, and so on.

There is no doubt that we have ability to *recognise* that we are being told this or that. It is only in such terms that we can hope to do justice to our reliability in telling when we are being told something despite the subtlety of the marks of telling. We acquire the ability through being initiated into the practice of knowledge-transmission. The fact that there is such a practice contributes to explaining how we all learn how to go about telling in similar ways, and how we learn to recognise tellings. In the next section I shall outline an account of practices that will be applied to the practice of knowledge transmission.

3. Practices

In the sense intended here, a practice is *an essentially rule-governed activity or group of inter-related activities*. Playing in a game of football, for instance, is participating in a practice. This activity is individuated by the rules—it is the activity of playing a game that is defined by those rules.⁸ So it is essential to the activity that it is governed by the rules: there could not be this activity but for those rules.

How should we think of the relationship of players to the rules? Obviously, players are subject to the rules in that the rules apply to them. The rules also set limits to what counts as playing the game. Even if, as in football, games can proceed although some rules are broken, there are limits to how anarchic play can be if a game is to proceed. But what should be said about how a player ought to behave or is committed to behaving in virtue of being a player? Some might be attracted by the idea that the rules do not generate any normative commitments or obligations beyond those incurred by aiming to play. If we ask what reason is there for players to follow the rules, the answer by this account would be, ‘Because they intend to play the game and could not do so without following the rules at least to some extent’. On this way of thinking the only normativity we need to acknowledge in connection with rule-governed activities is instrumental—following the rules by-and-large is what you have to do if you are to play. The rules have no normative force; it’s just that you won’t count as playing unless by-and-large you conform to them. If this is right there need be nothing wrong with a football player’s being in breach of a rule. Flouting a rule evidently need not be at odds with being a player and so does not put one in breach of the commitment incurred by intending to play the game—the commitment to doing what is necessary to that end. All we have is a clash between what the player has done and what some rule prescribes. I do not think that this is satisfactory. There is something wrong about being in breach of a rule and this is reflected in the fact that breaches of the rule are subject to legitimate criticism. What is it that makes the criticism apt? I suggest that to explain this we need the notion that it is in the nature of rules governing a practice that those who participate in the practice incur a (normative) commitment to following the rules. Behaviour that flouts a rule may legitimately be criticised because it is in breach of that commitment. Incurring this

⁸ Strictly speaking, any set of rules of football individuates a variant of the game. Games with somewhat different rules may be recognisable as variants of football.

commitment does not entail that you ought, straight out, to follow the rules. There can be games so appalling that no one ought to play them and therefore no one ought to follow their rules. By my account it remains true that if you play such a game you incur a commitment to following the rules. The commitment amounts to something like this: you ought to avoid becoming a player, or continuing as a player, and not follow the rules. There are two ways to discharge this commitment: one is to carry it out, that is, follow the rules, the other is to withdraw from the game, thereby removing the condition in virtue of which the commitment was incurred. It could be that you ought to withdraw from the game. Note that on this view the mere existence of a practice gives no one, not even participants, an obligation to follow its rules. The sense that practices do not generate obligations to follow their rules may explain the resistance that some have to the very idea that there is an intrinsic normative dimension to participation in practices. But we can acknowledge the point about obligation while still acknowledging that practices have an intrinsic normative dimension. What we need is a weaker normative notion than that of obligation—the notion of a commitment.⁹

The notion of a practice is philosophically interesting for three reasons. The first has to do with the metaphysics of human action. Certain activities are essentially-rule governed. People learn how to engage in those activities through being initiated into the practice and gaining an appreciation of what it demands. The second has to do with the explanation of action. The fact that there is a practice that commits participants to acting in a certain way figures in the explanation of why participants act in that way. Of course, this happens because people recognise what the practice commits them to doing. Nonetheless, the fact that the practice does commit them to acting in certain ways explains why they do act in those ways to the extent that they do, albeit via their recognition of what the rules prescribe. The third reason why the notion of a practice is philosophically interesting is epistemological. Observers and competent players of a game know how play is supposed to be shaped by the rules. This enables them to form reasonable expectations about the course of play. In a team game like football this in turn enables players to coordinate with others on their team and to be prepared for moves on the part of the opposing team that are likely or possible. A feature of practices that will turn out to be of

⁹ Commitments generated through participation in practices are structurally analogous to the commitments incurred by beliefs and intentions. On this, and on commitments generally, see Millar 2004. My thinking on such commitments is much influenced by John Broome's work. See, for instance, Broome 2002. However, his larger framework should not be read into mine.

importance for the present discussion is that knowledge of the practice can ground expectations about how participants are liable to act even in the absence of knowledge of the participants as individuals. This is not to deny that in some games, for instance, tennis, a player's individual quirks and preferences have an important bearing on what in detail he or she will do. The point is that even then you do not need to know the quirks or preferences to make out the general shape of what a player will do.

The sorts of practices I have mentioned thus far are ones that are instituted by design and have more or less well-defined rules by reference to which participants guide what they do. There can be practices for which rules are never formulated but which develop because certain ways of proceeding come to be seen as *the done thing* or *the way to go about things*. This is manifested by the fact that there arise mutual expectations of conformity, and a willingness to do the done thing and to view departures, whether by oneself or others, as open to legitimate criticism. What is it for there to be rules of such a practice? Consider, for instance, the practice of giving and receiving invitations. It's a rule of the practice that you don't issue invitations without planning, and implementing plans for, the activity in question. There is a certain pattern of behaviour that is in keeping with the rule: issuing invitations and planning, and implementing plans for, the activity in question. And there is behaviour that is not in keeping with the rule: issuing invitations and not planning, or not implementing plans, for the activity in question. To say that there is a rule along these lines is, roughly speaking, to say that behaviours of the former kind are acknowledged by those giving or receiving invitations as right or correct or appropriate, and behaviours of the latter kind are acknowledged to be wrong or incorrect or inappropriate. Part of what these acknowledgements involve is an understanding that one's behaviour is expected to conform to the first pattern and not to conform to the second. These informal practices, as one might call them, have epistemic and explanatory utility just as those that have been instituted by design. That there is a practice of giving and receiving invitations enables those who issue invitations to have a reasonable expectation that those invited will respond in a certain way, for instance, accept the invitation to dinner and then turn up at the set time. It enables those invited to have a reasonable expectation that the activity will indeed take place. All this may proceed smoothly with little detailed knowledge of the characteristics of those with whom one is interacting beyond the fact that they may be expected to *do the done thing* with respect to the giving and receiving of invitations.

4. *The practice of knowledge transmission*

Suppose you wish to make a table to a design that you have prepared. The design itself was selected by you in keeping with whatever needs or inclinations motivated you to embark on this table-building task. How you should proceed is determined by, among other things, the design, the characteristics of the wood you are using, the availability of tools, and the time at your disposal. Suppose you share a workshop with others, each of whom works on his or her own woodworking projects. If any of you is to know what others are doing you need to know from them what they are aiming to make. If you want to know how their work will proceed in detail you need to know from them the design to which they are working. If you want to know why they want to make what they are making in the way they are making it then, again, you have to ask them. Imagine now a small community of people who work in wood and who agree to make a range of products in specified styles. We now have a practice the rules of which include prescriptions for the products to be made and the style to which they are to be made. Without knowing anything about the inclinations or predilections of the others each can know quite a lot about what others are doing. They will be making one or other of the specified products in the specified style. To know why they are doing this it is enough to know that they are participants in this communal workshop. As such they have submitted themselves to the prescriptions in play. When people join the workshop they are initiated into the range of products and styles and learn how to execute them. Like the others they come to view themselves as doing the things they are doing because these are the things you do around here. They have a standing expectation that others will proceed in keeping with the practice. Appraisal and correction of what novices do will serve to keep the practice in play. Of course, there could be rebels who work on products or styles that are not specified, but they would not be rebels for long. Either they would be cast out or they would be seen as innovators and alter the range of specified products or styles.

The practices involved in communicating by means of a particular language are vastly more complex than any in the idealised workshop I have described. While those in the workshop opt into it, we do not opt into the linguistic practices associated with our native language. While those in workshop will represent the prescribed products and styles to themselves, we do not represent the rules that govern linguistic practices to ourselves. But in important respects these

latter practices have a role that is analogous to that of the workshop practice. (i) The fact that a practice is in play explains how others become initiated into it. Initiates are exposed to patterns in usage and to correction and thereby learn to use language for a variety of purposes. They are not like anthropologists making observations, trying to find plausible generalisations about how others behave and then trying to apply them. At the start they lack the resources for weighing up evidence and framing generalisations. Encounter with use of language shapes how they come to use language, but not by providing evidence-based justification for hypotheses about how others use language. (ii) The fact that there are practices into which we have been initiated, whereby words in our language are used for particular purposes, enables us to grasp the linguistic intentions of others who use our language. We do not approach them with an interpretative problem—what to make of the sounds they utter or the inscription they write. And we certainly do not project our own patterns of use onto them. We are already geared up to viewing them as fellow speakers of our language and to recognising what they are doing with the words they use: asking us this, telling us that, and so on.¹⁰ (iii) An (at least) inchoate sense that there are right and wrong ways to speak and write serves to reinforce our own continuing conformity to the demands of using our language. It is easy to overlook this because so much in our use of language and our understanding of uses of language is unreflective. It seems plausible however that our willingness to adjust our usage when we find that it is out of kilter with common practice, and we are not in the business of being innovative, testifies to our having a sense of correctness and incorrectness in linguistic matters.¹¹

We need to view the practice of knowledge transmission in the light of these considerations. Telling is a move in this practice. Because we have been initiated into the practice we know how to tell things, as opposed for instance to expressing personal opinions, or giving advice, and we are able to recognise when we are being told things. But above all we have an understanding that telling someone that *p* is the accepted way to get across to that person that *p*. The practice of knowledge transmission functions to facilitate the transmission of knowledge. This entails (a) that it actually does serve to facilitate the transmission of knowledge and (b) that it is sustained because of this. It is compatible with this that there are acts of telling that do not inform. I may

¹⁰ Mulhall 1990, especially in ch. 4, is instructive on the problems of explaining the understanding of utterances, on the model of interpreting sounds conceived as devoid of meaning.

¹¹ When corrected for misusing the term ‘arthritis’ Burge’s (1979) patient does not respond by saying that he what he means by ‘arthritis’ is different. He adjusts his usage.

tell you something that I know to be false. Or I may speak truthfully though the conditions obtaining are not apt for me to inform you. Suppose, for instance, that I tell you that dinner is ready. There are markers that this is an act of telling including the following: I say, ‘Dinner is ready’; what I say is something that can be told—it is the sort of thing that one person can tell another; my manner of speaking is apt for giving information—it is not as if I am trying to do something else; the context is not such as would indicate that I am, for instance, joking or mimicking someone else. Yet my telling you that dinner is ready may not be apt to inform you of this. It may be that you have no reason to think that dinner is in the offing and, because of that, do not know what to make of my utterance. A variant case would be one in which you have no reason to think that I know that dinner is ready. Imagine that there are people working away in the kitchen preparing dinner. You are in the sitting room. Having been away all day, I enter the house without having passed through the kitchen and, without further ado, tell you that dinner is ready. Again, you may well not know what to make of my utterance. It would not be apt for informing you without some explanation, because the circumstances are not such that you are in a position to know that I am in a position to inform you. If we think of telling as a move in a practice then these cases, along, with cases of deceptive telling, can be regarded as cases of infelicitous telling. What counts as felicitous telling depends on the function of the practice of telling. Since the point of the practice is to transmit knowledge, felicitous telling must, therefore, be (at least) apt to inform.

On the practice theory we know when we are being told things because we have a grasp of the practice of telling. The efficacy of telling as a means of informing turns on there being a shared understanding that telling is supposed to inform. If it were not the case that telling is supposed to inform we would need to provide or exhibit signals additional to merely telling if we were to get others to take us to be informing. We have no need of such additional signals just because tellings are understood as acts that are supposed to be informings.

5. Knowing through being told

The fact that there is a practice of knowledge transmission goes a considerable way towards explaining how it is that we so often acquire true beliefs from testimony. The following factors deserve emphasis.

- (a) The common currency of knowledge-transmission comprises easy-to-verify factual matters.¹² The sorts of things that informants routinely tell, with the expectation of being believed without further ado, are the sorts of thing that are readily knowable and which they could plausibly be taken to know. That, of course, is no guarantee that informants speak from knowledge. Its significance is that suspicions are naturally aroused if what an informant tells is not readily knowable or not the sort of thing the informant could readily be taken to know. It seems plausible that the practice of telling would not be sustained unless easily verifiable matters were routinely vouched for, received as knowledge, and subsequently confirmed.
- (b) The existence of the practice of knowledge transmission tends to constrain what people are likely to tell in such a way as to reduce the chances that what they tell is not from knowledge. The constraints are most forcefully present when informants are in social environments in which they are liable to run into those whom they inform or purport to inform. The practice constrains deception because in those environments deception is liable to become known by the deceived and others to the detriment of the deceiver and to the efficacy of the deceiver's subsequent communications, whether sincere or not. Similarly, the practice limits the chances of sincere telling not being from knowledge in the social environments in question. Tellings that are not from knowledge run the risk of letting recipients down because recipients come to know that they were misinformed or that their informant had no good basis for what was told. (Even if what was told is true, learning that the informant had no basis for what was told properly undermines our confidence in what was told so long as we lack independent means of knowing.) In a complex society there are many occasions when we should be circumspect about what people tell us. This is especially clear when they are trying to sell us things and when their gain from our believing them or, at least, not contesting what they tell us, outweighs any costs they would incur by misleading us. This sad fact does not undermine the general point that in social environments in which informants are liable to incur costs through not speaking from

¹² What counts as easy to verify is highly context dependent. In the context of a physics lab it might be easy to verify that a particular kind of particle has passed through a cloud chamber.

knowledge, we shall acquire true beliefs by accepting what we are told on many more occasions than we shall acquire false beliefs. To appreciate this we need only consider just how often we receive accurate information from members of our family, friends, colleagues, and people in organisations with whom we interact. (Where the latter are concerned it is easy to focus on the occasions when organizations misinform us, forgetting the many occasions when they inform us.)

It would be nice and neat if the above considerations sufficed to explain how we can routinely gain knowledge from testimony. Unfortunately they do not. The problem is that they do not explain why someone's telling us something can settle the matter. More specifically they do not explain how someone's telling us that p in the way that he or she does can indicate that p. The natural way to do this is by introducing the notion of trustworthiness. What indicates that an act of telling is truthful is that the speaker is trustworthy with respect to what he or she is telling. I take it that trustworthiness has two components, *sincerity* and *competence*.¹³

Sincerity, as I shall understand it has three strands: you are sincere in telling me that p if and only if (i) you believe that p, (ii) you believe that you speak from knowledge, and (iii) you intend that I should come to know that p through telling me. Condition (i) needs no explanation. The rationale of condition (ii) is provided by the nature of telling as an act by which you give it to be understood that you know. If you believe what you say but think that your standing on the matter is less than knowledge then you are misleading those you tell. The rationale for condition (iii) is similar. Suppose that I tell someone that her husband, my friend, is not having an affair. He is not and I know this, but I aim to get her to believe, falsely, that her husband is having an affair, assuming that she will think that I am covering up for my friend. Since telling her that her husband is not having an affair is an act by which I give her to understand that I intend that she should come to know this, I would mislead her if I lack any such intention.

Competence with respect to whether p is knowing whether p, which is to say, knowing that p if p. A crucial dimension of competence is being in command of the requisite ways of telling. My son meets this condition because he knows the title of the book about which I am enquiring and can determine by looking whether a book of that title is in my study. Obviously, merely

being in command of the requisite ways of telling is not enough. The ways of telling have to have been exercised. My son meets this condition because he has seen that the book is on my study desk.

If you are trustworthy on a range of matters then you are akin to a reliable instrument on these matters.¹⁴ Since you are competent with respect whether something in the range holds true you are highly unlikely to think you know something in the range when you do not. Since you are sincere with respect to matters in the range you are highly unlikely deliberately to mislead others on these matters in any of the ways identified above. So on these matters it is highly unlikely that you will tell people that *p* when not *p*. However, the fact that you are trustworthy obviously does not by itself explain what furnishes *me* with a way of knowing on a specific occasion that your telling me this or that is truthful. I need to know that you are trustworthy with respect to what you are telling me. We need to look into how experience can put us in command of such a way of knowing. There are issues here that are important for epistemology generally and not just for the understanding of knowledge from testimony.

6. Experience and the acquisition of ways of knowing

What might be called the standard model of empirical knowledge envisages that experience furnishes us with facts about particulars, which provide the basis for generalisations, which can then be applied in explanation and prediction. For instance, to know that the fact that a is F indicates that a is G one needs empirical support for a covering generalisation that connects a's being F with its being G. This might be that everything F is G. Obviously, there can be support for such generalisations, and one can be justified in thinking them true or at least in placing some fair degree of confidence in their being true. On this model it is our stance on the generalisation that provides us with a way of telling that something is G from the fact that it is F. But there is an alternative way to conceive of how experience furnishes us with ways of knowing. We may think of experience as forming in us an ability to recognise ways in which certain situations are

¹³ Elizabeth Fricker (1994) thinks of trustworthiness has having these two components. Much of what I say is in the spirit of her account though, unlike her, I am not working within the framework of a conception of knowledge as built out of justified true belief.

¹⁴ Links between knowledge from testimony and knowledge via instrument reading are suggestively explored in Sosa (2006).

significant. By ‘experience’ here I include our perceiving that this or that is so but also the experience by which we are bombarded with observations, corrections, explanations, and so on by those who surround us. Through such experience we get the hang of telling that it has been raining from the wetness of the streets and pavements, of telling that someone wishes to enter from the knocking sound at the door, of telling that it has been very cold from the frost on the grass, of telling that cattle are around from the sound of mooing, and so on. When I say that we get the hang of these things I mean simply that we learn to do them. The upshot of the learning process is a recognitional ability, in particular, an ability to recognise a situation as having a certain significance. We should think of the role of experience accordingly. Its role is in some respects akin to the role of experience in shaping an ability to saw wood straight, or cultivate orchids, or cook fish perfectly. In these latter cases we take steps towards a certain end, and gain feedback from our faltering attempts, perhaps assisted by guidance from those who know about these things. The upshot of the learning process is mastery of a technique. The role of experience is that of honing that technique by trial and error, so that in time it becomes an ability to do something well. The test for whether we have achieved this is whether we have acquired the technique—whether we saw well, cultivate orchids well, cook fish well, as the case may be. Now consider how, without guidance from others, someone might learn that the presence of certain droppings indicates that rabbits have passed by. (Imagine someone stranded on a deserted island.) This will involve observing rabbits produce such droppings and learning to discriminate the rabbit droppings from others that might be around. Initial steps towards this may be faltering. Perhaps there are droppings very like rabbit droppings, produced by another small mammal. But in time the learner catches on to recognising just the right sort. The upshot again is a certain technique, in this case judging that rabbits have passed on the basis of the droppings. Getting the technique right is acquiring a recognitional ability—the ability to recognise the droppings to have a certain significance: that rabbits have passed by. The test for whether the ability has been acquired is whether over a wide range of cases in which an attempt is made the learner judges correctly that rabbits have passed by with a very high degree of reliability. The focus of the learning process is on identifying the right sort of dropping. The same applies to knocks at the door. Not any sound at the door is the right kind of sound to count as a knock. One has to learn to

discriminate the knocks from the bumps and scrapings. Learning to connect the sound with the knocks is just learning to recognise the right kind of sounds as knocks.¹⁵

The perceptual-recognitional ability to tell by looking that it's a robin on the bird table is exercised in telling, which is to say, coming to know that it's a robin on the bird table. The exercise of the ability is nothing less than the acquisition of knowledge. The very notion of the exercise of this ability is a success notion. Of course, it is possible that someone who has this ability should, on rare occasions, judge falsely while aiming to tell whether a bird is a robin. That just means that on those occasions the subject does not succeed in telling.¹⁶ Analogous considerations apply to the ability to recognise indicator facts as having significance. The ability to tell from the fact that there is a knocking sound at the door that someone wishes to enter is exercised when one does so tell. Having the ability is compatible with occasionally misidentifying a sound as being a knock at the door.

One could have the ability to tell that it has been raining from the wetness of streets and pavements with respect to an environment in which a certain sort of wetness on streets and pavements is a highly reliable indicator that it has recently rained. Suppose that around here this condition is met. There might well be other environments in which that sort of wetness occurs but is regularly caused by water sprayed from a truck to keep dust down. If that is so then with respect to those environments one could not in the same way tell that it has been raining just from the relevant sort wetness of the streets and pavements.¹⁷ But the fact that there might well be such environments does not falsify the attribution of an ability to tell to that it has rained from the wetness of the streets and pavements around here. It has a bearing though on what we should say about generalisations in relation to the kind of recognitional ability of which I am speaking. If there might well be such environments then any law-like generalisation to the effect that usually when streets and pavements are wet, in what we count as the rain-indicating way, would

¹⁵ The importance of discrimination for perceptual knowledge is emphasised in Goldman 1976. Here I am emphasising its importance for knowledge of what situations indicate.

¹⁶ The exercise of the ability is thus not identical with manifestation of a disposition that reliably yields true judgements. Such a disposition can be manifested in making a false judgement. In that case the ability is not exercised.

¹⁷ This illustrates the analogue of a fake barn case (Goldman 1976) for knowledge that a situation has a certain significance. Note that on my view there remains a difference between perceptual-recognitional knowledge, like the knowledge by looking that a bird is a robin, and knowing one thing from another, for instance, knowing that deer have passed from the tracks on the path that one sees. In the former case, the knowledge depends on one seeing the robin; in the latter case, the knowledge does not depend on seeing the deer.

be false, if they are understood as having a wide domain of quantification extending to wherever you find streets and pavements. So knowledge that it has rained recently from the wetness of the streets and pavements does not depend on knowledge of any such generalisation and we should not assume that any such generalisation need enter the thinking of those who can tell that it has rained in this way. The thinking of those with the ability in question is likely to be such that they would readily assent to the statement that you can tell from the wetness of the streets and pavements that it has rained recently. The truth conditions of this statement are sensitive to the context of utterance. Whether it is true or not depends on the environment's being favourable.

It is important to stress that those who have the abilities I am describing do not exercise them blindly. In exercising these abilities they do not infer with the help of a generalisation but they are not blind to the basis of concluding as they do. In the rain example they will have a perspective on which one can tell *from the wetness of streets and pavements* that it has rained recently. So we are not faced with a theoretical choice between a view on which subjects behave like calculators blindly churning out the right answer and a position that requires subjects to apply, and know that they are applying, a generalisation for which they have empirical support. When people exercise the abilities of which I am speaking they will at least typically know what they are doing. They will know that they are telling that it is raining from the wetness of the streets and pavements.

Returning now to the main theme, the suggestion I want to explore is that acquaintance with a person can furnish us with a sensibility whereby we can know of this person that he or she is trustworthy on this and that occasion.¹⁸ We are to think of experience as shaping a recognitional ability by which we can tell that *p* from *X*'s telling us that *p*. Talk of a sensibility is no mere verbal flourish. The view accommodates

- (a) the unspecifiability of the features of the person telling, and of the telling, and the context of telling, that make the telling an indicator of the truth of what is told,
- (b) the consequent unavailability of generalisations adequate to explaining how we know that the telling is truthful, and

¹⁸ The suggestion is in keeping with Miranda Fricker's (2003) emphasis on the importance of a testimonial sensibility.

(c) the centrality for such knowledge of recognitional abilities.

It is not a mystery that I come to know that the book about which I am enquiring is on my study desk at home from my son's telling me that it is. It is because I know him well enough to have acquired an ability to tell from certain of his tellings that what he is telling me is true. The tellings with respect to which this is so are identified by manner of speaking, content, context and backgrounds circumstances. Stepping back from the occasion in question I can identify more specifically some of the relevant factors. My son's competence is not at issue. Since I gave him the book's title, all he needed to do was look in my study. I'm not likely to be in the mood for fooling around, and he will know that. There is a dull matter-of-factness about the way he speaks, that contrasts with the sprightly tone he adopts when, to show up by gullibility, he tells me that there is a smudge on my nose. Even if he had been fooling around he would quickly put the matter right. The content of the telling is not the sort of thing he could have any interest in lying about, even leaving aside that I know that he would not lie on such a matter. These are merely pointers to features of the situation to which I respond. Ask me to spell out what the matter-of-factness amounts to and I could not do so. It would be futile to try to list these factors and invoke a generalisation to connect them with the truth of what my son says. The point is that I know him well enough to have acquired an ability to tell when he is being sincere. Indeed, it is constitutive of my knowing him in the way I do that I have such an ability.

7. An objection briefly considered

The suggested account is likely to provoke an objection from the consideration that often there can be knowledge from testimony when the recipient does not know the informant well enough to have the kind of discriminative ability illustrated in the case of my son's testimony.

Consider, for instance, informants like ticket clerks, shop assistants, and representatives of organizations who inform us about their products. We often do not know these people personally though we do acquire knowledge from them. It might be thought that there is nothing about them from which we can tell that they are trustworthy. Such a view overlooks the function of roles. Such people occupy information-disseminating roles and thus participate in the practices that

define the role. They have little incentive to mislead and the risks to their doing so are significant. They will have acquired training so that they will generally speak from knowledge on a range of subject matters. Providing that we are sufficiently discriminating about the roles they occupy and about the contents of what they tell us we can be in a position to know some of the things they tell us. Experience hones our discriminative powers.

What about the person in the street whom one asks for directions to the castle? Can we not obtain knowledge from such a person while not knowing him or her personally and knowing nothing of any information-disseminating role the person occupies? Maybe. It depends on how the scenario is spelled out and, in particular, whether the informant can be seen to be trustworthy. What matters for my purposes is that what we gain from such a person will often be less than knowledge. The communication may be useful nonetheless. Provided one exercises a modicum of discrimination it is entirely reasonable to act on the information we receive, for example, going in the direction given, when little is at stake or when one has no alternatives. But acting on information supplied does not require belief. All it requires is that there should be a reasonable chance that what one is being told is true. That will be guaranteed by the fact that there is a practice of knowledge transmission, along with the fact that there is nothing suspicious about one's informant.¹⁹

A rather different line of objection focuses on early learning. When we are young we take on board a great deal of information from parents, teachers and others. We are in no position to have acquired the kind of recognitional abilities that are central to my account. If early learning is to be conceived as the acquisition knowledge through being told then the knowledge will not meet the conditions I have laid down. Though it is widely believed that an epistemology of testimony must conceive of early learning in this way, my own view is that this is a mistake. We do acquire a picture of the world through early learning but the status of what we thereby acquire is not knowledge-through-being-told in the sense of knowledge from being told by this or that person, as I have conceived it. Knowledge of the latter sort retains its status as knowledge of this sort through recollection of the telling. It is like the knowledge that I have that my book is at home from my son's telling me this. We forget virtually all of the occasions of telling through which have imbibed the elements of our world picture. Much of what we imbibe is not knowledge at all. Some of it, for instance, basic geographical and historical facts of it becomes knowledge

through repeated encounters with the facts in question.²⁰ (It seems plausible that the status of knowledge of such basic facts is not a matter of having justification for the corresponding beliefs deriving from our encounter with the facts, but of having derived from those encounters an ability to recall the facts. Undoubtedly our picture of the world is reflected in what I am calling ways of telling. But I suggest that the epistemic role of elements of the picture should be understood in terms of their contribution to ways of telling. If we reverse the order, and aim to understand the status of ways of telling in terms of elements of a picture that seems to hang in the air, we court disappointment.

References

- Adler, J. (1994) 'Testimony, Trust, Knowing', *Journal of Philosophy* 91, 264-75.
- Austin, J. L. (1962) *Sense and Sensibilia* (Oxford: Clarendon Press).
- Broome, J. (2002) 'Practical Reasoning' in J. L. Bermúdez and Alan Millar (eds.) *Reason and Nature: Essays in the Theory of Rationality* (Oxford: Clarendon Press).
- Burge, T. (1979) 'Individualism and the Mental' in P. A. French, T. E. Uehling, Jr., and H. K. Wettstein (eds.) *Midwest Studies in Philosophy*, iv, (Minneapolis: University of Minnesota Press, 73-121.
- Burge, T. (1993) 'Content Preservation', *Philosophical Review*, 102, 457-88.
- Coady, C.A. J. (1992) *Testimony: A Philosophical Study* (Oxford: Clarendon Press)
- Fricker, E. (1994) 'Against Gullibility' in B. K. Matilal and A. Chakrabarti (eds.) *Knowing from Words* (: Kluwer), 125-61).
- Fricker, E. (1995) 'Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony', *Mind*, 104, 393-411.
- Fricker, M. (2003) 'Epistemic Injustice and a Role for Virtue in the Politics of Knowing', *Metaphilosophy*, 34, 154-73.
- Fumerton, R. (2006) 'The Epistemic Role of Testimony: Internalist and Externalist Perspectives' in Lackey and Sosa (eds.) (2006), 77-92.

¹⁹ This may be a source of the default-belief view.

²⁰ Compare Adler 1994, who makes the parallel point for justified belief.

- Goldman, A. (1976) 'Discrimination and Perceptual Knowledge', *Journal of Philosophy*, 73, 771-91.
- Lackey, J and Sosa E. (eds.) (2006) *The Epistemology of Testimony* (Oxford: Clarendon Press).
- McDowell, J. (1981) 'Anti-Realism and the Epistemology of Understanding' in Herman Parret and Jacques Bouveresse (eds.) *Meaning and Understanding* (Berlin: Walter de Gruyter). Reprinted in McDowell 1998.
- McDowell, J. (1982) 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy*, 68, 455-79. Reprinted in McDowell 1998.
- McDowell, J. (1998) *Meaning, Knowledge, and Reality* (Cambridge, MA.: Harvard University Press)
- Millar, A. (2001) 'The Scope of Perceptual Knowledge' *Philosophy*, 75, 73-88.
- Millar, A. (2004) *Understanding People: Normativity and Rationalizing Explanation* (Oxford: Clarendon Press).
- Millar, A. (2007) 'What the Disjunctivist is Right About', *Philosophy and Phenomenological Research*, 74, 176-98.
- Millar, A. (2008) 'Perceptual-Recognitional Abilities and Perceptual Knowledge' in Adrian Haddock and Fiona Macpherson (eds.) *Disjunctivism: Perception, Action, Knowledge* (Oxford: Clarendon Press).
- Moran, R. (2006) 'Getting Told and Being Believed' in Lackey and Sosa (eds.) (2006), 272-306.
- Mulhall, S. (1990) *On Being in the World: Wittgenstein and Heidegger on Seeing Aspects* (London: Routledge).
- Sosa, E. (2006) 'Knowledge: Instrumental and Testimonial' in Lackey and Sosa (eds.) 2006, 116-23.
- Welbourne, M. (1986) *The Community of Knowledge* (Aberdeen: Aberdeen University Press).
- Welbourne, M. (2001) *Knowledge* (Chesham, Bucks: Acumen)