

The Epistemological Significance of Practices

Alan Millar

1

Much of our knowledge concerns the thoughts and actions of others. We often know what people think and why they think it, and what people are doing and why they are doing it. Where we lack knowledge in such matters we might be able to form well-grounded beliefs as to what it is likely to be so or might well be so. I might understand why a colleague believes that a speaker might not be able to turn up when I hear that this colleague has just learned that the speaker is delayed at an airport. I might understand why a friend is going London when I learn that she is going to visit an interesting exhibition. In these cases we have an explanation of some attitude or action in the light of other attitudes the agent has. Unsurprisingly, on the basis of known attitudes we can often make reasonable predictions about what people will think or do in this or that circumstance. Understandings and predictions of this sort are routine because we routinely acquire from people information about what they know or believe and about what they desire or care about or intend. They simply tell us something about such matters.

Crucially, whether explaining or predicting we rely on assumptions about what it *makes sense or would make sense* for people to think or do in the light of their attitudes. But often we find people's thought or action intelligible, or correctly anticipate what will they think or do, despite having little if any information about their attitudes other than what we can glean from their situation and non-verbal behaviour. Watching a football match you see a player in one team trying to avoid being too close to two players in the opposing team. You understand that the first player is being marked by the others and wants to avoid being prevented by them from receiving the ball. This is despite your having no specific information about any of the propositional attitudes of the players beyond what you may infer about them from their circumstances and their non-verbal behaviour. Other examples easily reinforce the point. Suppose you receive an invitation to give a talk at some academic institution. All you receive is a brief email telling you about the nature of an event that is being arranged and asking you to give a presentation. You reply saying that you would be happy to come. You can readily understand that the email is an invitation. You expect that the sender will ensure that the

event takes place as described. You also expect him or her to take from your reply that you will come as arranged. Your response is reasonable though made in the absence of evidence concerning the personal traits, or particular propositional attitudes, of your interlocutor, beyond what you can gather from the email. What is remarkable about this is simply the apparent tension between the paucity of information we have about people's attitudes and the fact that we still seem to be able to arrive at correct understandings and predictions their behaviour. These understandings and predictions, it seems, have to be in the light of their attitudes yet we seem to know so little about their attitudes.

One response to the tension is to suppose that we work with a theory—a body of generalizations—that we are justified in accepting and that we apply to specific circumstances.¹ In relation to the case of the invitation the idea would be that you work with, among others, the generalization that when people at academic institutions issue invitations they are likely to ensure that the event in question will take place. No doubt something along these lines is a presupposition of your thought and action in response to the invitation. But how plausible is it that you have evidence adequate to justify your acceptance of this generalization? Maybe you have had loads of invitations and everything has worked smoothly. Maybe colleagues have told you about invitations they have received and things have always worked smoothly. But maybe this is your first invitation and you have very little experience bearing on the matter. Yet you will no doubt be pretty confident that the event in question will take place and with good reason.

Another reason for thinking that the suggested answer is inadequate is that it is far from obvious that we always have suitable generalizations available. Someone utters a sentence in our presence. We understand what is going on because, among other things, we identify the person's utterance as one of *telling* us that something is so, as opposed, say, to voicing an opinion or giving us advice. What generalizations are at our disposal in a case like this? We respond to a host of cues that we would be hard put to specify except in the most general terms. No doubt the content of the utterance, the manner of speaking, and the situational context, are relevant, along with much else besides. But how are we to spell out, for instance, just what kind of manner of speaking counts? It seems plausible that we do not actually take in at the level of belief or judgement which features of the utterance mark it as a case of telling. As in the cases in which we recognize a familiar face we respond to a total *Gestalt*.

¹ For this view, see Fodor 1987: ch. 1, which includes discussion of an example involving an invitation at pp. 3-4.

Our response is recognitional rather than based on assumptions about features. But if we do not take in at the level of belief or judgement which relevant features are present we can hardly apply an appropriate covering generalization.²

Faced with such considerations there might be some temptation to respond by saying that in cases where we genuinely understand what is going on, rather than merely guess or conjecture, we must work with a body of generalizations that we are justified in accepting, because there is no plausible alternative. But there are alternatives.

The alternative I consider here invokes the notion of a practice, conceived as an essentially rule-governed activity.³ Practices, I claim, have enormous epistemological significance in relation to understanding and predicting human behaviour. It is largely thanks to their existence that we are so well able to understand and anticipate a great deal of what people think or do despite the paucity of information that we have about them. We rely on reasonable assumptions about what they are supposed to be doing in view of the activity in which they are engaged—an activity for which certain ways of proceeding are prescribed and others proscribed. In addition we presume, in the absence of countervailing considerations, that, by and large, people engaged in certain sorts of activities—the sort that the theorist identifies as practices—do by and large what they are supposed to in that role. So there is a high-level generalization in the picture—one that forges a connection between what people are supposed to do and what they are likely or liable to do. And there are propositional attitudes in the picture—those that a participant in a practice can reasonably be expected to have. I shall return at the end to the status of the generalization.

Consider a somewhat idealized example designed to highlight the contrast between a situation in which there is no practice and counterpart situation in which there is.⁴ Imagine a group of artistic painters. Although they share a workshop, each works individually on whatever it is that he or she is painting. Nothing other than their inclinations and skill dictates how they proceed. The workshop is simply a shared space for pursuing individual projects, rather than cooperative projects. If any painter is to know what others are painting they have to look and see. At the early stages of a painting they will have to ask the artists what they are doing. That is to say, they will need to discern the propositional attitudes in the light of which

² I explore the role of generalizations further in my contribution to Pritchard, Millar and Haddock 2010.

³ I discuss simulation theory, another alternative, in Millar 2004: ch. 8.

⁴ I use an analogous example for similar purposes in Millar 2010.

the activities of the other artists make sense. Now, contrast this situation with one in which the various painters are engaged in a cooperative enterprise. Each has an allotted role in the production of a painting based on a drawing provided by a master painter, and each has a grasp of the roles that others have in the production of the painting. The upshot is that the production of the painting is an essentially rule-governed activity—a practice. The rules are prescriptions for producing some portion of the work. The practice is the production of the painting according to those rules. To be a participant in the practice to engage in it and thereby incur a commitment to proceeding in accordance with its rules, including whatever rules govern the role to which one has been assigned. Let us suppose that it also involves having some grasp of the roles assigned to others since there will sometimes need to be communication between those assigned to one role and those assigned to others. In these circumstances each knows what he or she is supposed to be doing (in other words committed to doing). Each has reason to do what he or she is supposed to be doing since there is no point simply because each is a participant in the collective activity. Each has some understanding of what others are supposed to be doing and, since others are participants, may justifiably presume that they will do what they are supposed to do. For this purpose all that each has to know about the others is that they are participants and have such-and-such a role. They do not need to know about them or their artistic predilections, though they will presuppose that participants have attitudes appropriate to their roles.

In large institutions like universities or firms or armies, people have to interact with others whom they do not know personally. So it just as well that institutions have organizational structures involving various roles and offices. A Dean of Faculty, for instance, occupies an office defined by certain duties. This amounts to being a participant in a practice that consists in doing the things Deans are supposed to do in virtue of their office. If members of institutions are to know what is to be expected of certain people whom they do not know personally they have to know which role or office they occupy and that means that there have to be indicators of roles or offices. In all institutions role titles serve this purpose. In armies uniforms indicate role titles. Members of an institution need to know the significance of these indicators and what is to be expected of those with whom they have to interact in the light of the roles or offices that they occupy. On this basis they can form reasonable expectations of how people will act and form reasonable understandings of why they have acted in this or that way in the light of what the role or office demands that they do. Where an action is one of a range of actions open to a person occupying a role, knowing the role and what it involves

will not tell you why the person acted as he or she did. Nonetheless, will tell you much about the responsibilities that structure what the person does.

In the examples I have been considering practices implicate the rules are explicitly instituted and which participants can find out about. Not all practices are governed by rules that participants represent to themselves. Arguably, there is a practice of issuing and responding to invitations to give academic presentations. Among the rules would be: issue an invitation only if you are making arrangements for the event to which you are inviting people; accept an invitation only if you are prepared to do what you are invited to do. These rules spell out what participants in the practice would acknowledge as *the done thing* in the way of, respectively, issuing and accepting invitations.

The general idea that I shall develop here is as follows:

- (i) that we often know or make reasonable judgements about why others think or act as they do or about what they are likely to, or might well do, in some circumstance, based on assumptions about what they are supposed to do given the activity in which they are engaged, where that activity is of the sort that the theorist identifies as a practice;
- (ii) that our understanding of what people are supposed to do when so engaged derives from our grasp of appropriate and inappropriate ways of proceeding in relation to the activity in question; and
- (iii) that this mode of understanding depends on its being true, and on our presuming, that by and large participants will think or act as they are supposed to in virtue of the commitments they incur as participants in those practices.⁵

In sections 2 and 3 I focus on how we should conceive of the normative dimension of practices. In section 4 I consider how to make sense of there being practices that lack explicitly instituted rules. In sections 5 and 6 I illustrate the general idea captured by (i) to (iii) with reference to linguistic practices.

⁵ I am building on a conception of practices set out in Millar 2004.

It is worth paying close attention to how best to characterize the normative dimension of practices. As well as being of some interest for its own sake this is necessary in order to address any doubts there might be about whether there are practices as I conceive of them—that is, as essentially rule-governed activities.

For the sake of simplicity let us reflect a little further on the example of football. The game is in a sense defined by the rules. Necessarily, to play in a game—to be a participant in the practice of playing a game of football—is to be subject to those particular rules. If you are not subject to those rules then it could not be that game you are playing—not that activity in which you are engaging. But merely being subject to the rules does not suffice to characterize how a player stands to the rules. Walking in a public park I might be subject to a rule forbidding the dropping of litter—a rule that is imposed by the local authority. My behaviour admits of evaluation in terms of the rule: it might comply with the rule or be in breach of it. But being subject to the rule might have no influence whatever on my behaviour. I might walk in the public park but be ignorant of, or indifferent, to what the rule requires. But a rule can *govern* one's behaviour in a sense that entails that the rule, or at least what the rule requires of one in particular circumstances, guides one's behaviour. One way in which a rule can guide one's behaviour is through functioning as a prescription that one represents to oneself and aspires to follow. For instance, in a case in which I know that there is a rule forbidding the dropping of litter I might try to ensure that I comply with it. The rule, via my knowledge of it, plays a role in explaining my behaviour. A condition of that sort is not met in all cases of governance by a rule. There are cases in which one is in a sense guided by a rule even though one does not represent the rule to oneself. These are cases in which one has learned that a certain way of proceeding is prescribed or proscribed. I shall shortly say a little about these.

There is, I claim, a normative dimension to being a participant in a practice that is captured by saying that participants incur a (normative) commitment to following—being governed by—the rules of the practice. A player in a game of football incurs a commitment to following the rules of football, and thus to avoiding doing certain things—the sorts of things that are breaches of the rules. Commitments incurred through participation in practices play a role both in explaining why people act as they do, and in enabling others to understand and anticipate their actions. That is because the particular ways in which actions make sense, or would make sense if done, can depend on their being actions that the agents are committed

to undertaking. We know why Federer will serve at a certain point in the game. It's because he is playing a game of tennis and the rules demand that he serve at this point. As we might say, he has to serve just then. That amounts to saying that he is committed to doing so. He is committed to doing so because he is committed to following the rules and the rules demand that he serve.

To make clear what all this amounts to I need to say something about normative commitments.

3

The operative conception of the commitments incurred by participating in a practice is modelled on a conception of the commitments incurred by beliefs and intentions.⁶ The general structure of those commitments is as follows:

- (1) There is a commitment-incurring condition.
- (2) There is something that the commitment-incurring condition commits one to doing.
- (3) The commitment amounts to its being the case that one ought to avoid remaining in the commitment-incurring condition while not doing that which it commits one to doing.

Schematically, the commitment-incurring condition might be *my believing that p and believing that if p then q*. As we might ordinarily put it, my believing these things commits me to believing that q. If (3) is correct that cannot be exactly right, though it is near enough right for ordinary purposes. If (3) is correct then the commitment, as just understood, amounts to it being the case that I ought to avoid believing that p and believing that if p then q while not believing that q. That is too strong. I cannot plausibly be taken to be in breach of any commitment just because I lack a belief in something entailed by something I believe. The assumption here is (i) that it must be possible to comply with commitments one incurs, and (ii) that since infinitely many propositions are entailed by what I believe it is impossible to comply with the supposed commitment. The commitment we are looking for must be to

⁶ The latter commitments should not themselves be viewed as grounded in practices. As I see it, they are grounded in, respectively, what it is to believe something and what it is to intend to do something.

believing that q if one gives a verdict on whether q . I shall be in breach of the commitment if I believe that it is not the case that q , while retaining the commitment-incurring beliefs. But there is a sense in which withholding judgement on whether q is a verdict on whether q , by contrast with merely not believing that q and not believing that it is not the case that q . So I would be in breach of the commitment if while retaining my commitment-incurring beliefs I withhold judgement on whether q . There does seem to be something amiss about being in this state. It would be one in which I have a certain picture of the world on which p , if p then q , and q , yet not realise that on this picture it is true that q .

One might think that the commitment just referred to derives from a requirement of rationality: do not retain the belief that p and that if p then q while giving a verdict on whether p other than belief. If a rationality requirement is something such that if you fail to satisfy it you are in some way irrational then there is a problem. For one can fail to satisfy the supposed requirement through ignorance of an entailment of what one believes that has nothing to do with irrationality. Even in the schematic case considered above one might fail to put together the information that p and that if p then q and not see that it entails that q . Ignorance of the entailment need not be a case of irrationality. Being in breach of a commitment of the sort under consideration is not necessarily a breach of a rationality requirement. Nor is it the case that discharging a commitment is necessarily an improvement or enhancement of rationality. Reverting again to the schematic example, I can discharge the commitment in question either by believing that q or by ceasing to believe that p or ceasing to believe that if p then q . (Notice that this would not be so if the commitment amounted to a requirement to believe that q .) The commitment is silent on what is the best way to discharge it. I might irrationally believe that q , thereby discharging the commitment, but with no enhancement of my position vis-à-vis requirements of rationality. But commitments of this sort do have something to do with rationality. I suggest that we think of there being *an ideal of rationality* such that one ought not to give a verdict other than belief on anything entailed by what we believe. The sense of 'ought to' in this context is that of 'ideally would'. What rationality requires of us is that we respect this ideal in our thinking. So if, for instance, we find that we are in the position of believing certain things, knowing that these things entail that q but are unwilling to accept that q , then we ought to do something about our condition. In particular, we ought to seek the best way to discharge the commitment incurred. In this context the sense of 'ought to' has the force of a requirement such that breach of it is a failure of rationality. Closely analogous points apply to commitments incurred by intentions—specifically, commitments to taking the means necessary to carry out our intentions. If I

intend to ϕ , and if ψ ing is necessary if I am to ϕ , then I incur a commitment to ψ ing. That amounts to it being the case that I ought to (ideally would) avoid retaining the intention to ϕ while never getting around to ψ ing. I can fail to discharge the commitment due to ignorance or through not succeeding in ψ ing, without having changed my mind about ϕ ing. (For an instance of the latter, suppose I intend to reach the gallery before it closes at 4pm, run to get there on time but arrive just as the doors close. I have not had time to change my mind and ensure that I discharge the relevant commitment.) So, for reasons analogous to those explored above, incurring the commitment is not the same as being subject to a rationality requirement to the effect that one ought (in the requirement sense) to avoid retaining the intention to ϕ while never getting around to ψ ing. But plausibly there is an ideal which rationality demands that we respect. In the requirement sense we ought to seek to avoid retaining intentions while not taking the necessary means.

The commitments incurred by participation in a practice have a closely analogous structure to those incurred by beliefs and intentions. Participating in a practice, G , incurs a commitment to following the rules of G . To a first approximation that amounts to it being the case that one ought to avoid continuing to participate in G while not following the rules of G . The 'ought' here does seem to be the 'ought' of requirement, the idea being that non-compliance is necessarily some kind of fault. (Though see below for an important qualification.) As before there are two ways to discharge the commitment, by ceasing to participate in the practice and by complying with the rules.

In the case of the commitments incurred by beliefs and intentions it is clear that there is reason to avoid representing the commitments as obligations or requirements to do the thing to which one is committed. It might be wondered whether there is an analogous reason to avoid representing the commitment incurred by participation in a practice in terms of an obligation to follow the rules. Suppose that I am playing a game of football. Is there not a straightforward sense in which I ought to follow the rules given that I am a player? There might be but we are considering a commitment that is conceived as arising simply from the fact that one is player, without regard to any contractual or moral obligations that one might have through being a member of a team. It is reasonable to ask, 'Why should the mere fact that one is a player and that there are such-and-such rules yield any conclusions about how I ought to behave?'. Many years ago Philippa Foot posed a related challenge in connection with etiquette. She said

... although people give as their reason for doing something the fact that it is required by etiquette, we do not take this consideration as *in itself giving us a reason to act*.

Considerations of etiquette do not have any automatic reason-giving force, and a man might be right if he denied that he had reason to do 'what's done'. (Foot 1972/1978: 161).

This can be turned into a problem for the theory of practices since etiquette can plausibly be taken to be a cluster of rules governing such things as issuing and responding to invitations, behaviour while dining, and so forth.⁷ Expanded in line with this, Foot's point is that though the rules of etiquette might prescribe that such-and-such be done, this by itself does not give anyone a reason to do that thing. Actually, this seems to me to put the point too strongly. For there does seem to be a sense in which participating in the practices of etiquette gives one a reason to do 'the done thing'. Arguably, having as a participant incurred a commitment to doing the done thing favours doing 'the done thing' at least to the extent of conferring a point on doing so: it discharges the commitment. But even if this is right, the key point that Foot is after stands: to anyone considering whether it matters whether one behaves in keeping with rules of etiquette the mere fact that there are these rules does not show that the right thing to do is to comply. I suspect that this kind of consideration is one reason why some are sceptical that there are practices that have an intrinsic normative dimension. If so it is not a good reason for scepticism since the idea that practices have an intrinsic normative dimension can accommodate Foot's key point. The normative dimension does not consist in it being the case that a participant is required or obliged to comply with the rules. It consists, to a first approximation, in it being the case that one ought to avoid continuing to participate in the practice while not complying with the rules. In the etiquette case participating would simply be a matter of a being in a social milieu governed by rules of etiquette. There might be good reason to avoid staying in that milieu.

Notwithstanding those considerations some might be sceptical that there are such things as practices as I am conceiving them. Here is a line of thought that might be advanced at this point.

⁷ This might seem to be at odds with my claim that practices are *essentially* rule-governed activities. It is possible to set dining tables without being governed by, say, rules of etiquette that prevail in Buckingham Palace. Obviously, one can set tables and not be so governed. But where those rules of etiquette prevail there is a sense in which a table is not properly set if its layout does not comply with the rules. To be set is to be laid out as the rules prescribe. The underlying general point is that the activity in question is not just that of, for instance, setting tables for dining in some way or other, but setting tables for dining in keeping with a certain set of rules.

To understand playing games that are defined by rules we just need to understand instrumental rationality. A player of a game of football intends to play the game and so is committed what it takes to play the game. To play in the game the player has to comply with the rules at least to an extent required to avoid being sent off. So players incur means-end commitments to complying with rules to that extent. The rules simply set out the modes of behaviour such that if you do not accord with those modes of behaviour you incur a sanction.

The suggested account seems plausible initially because it focuses on a particular type of example—an activity that people engage in intentionally with rules that they represent to themselves. Participants in such activities can view following the rules to some extent at least as a means to whatever further end they have in playing the game or just to the end of playing the game. But this model does not fit all practices. Arguably, there are practices the participants in which need not view participation as an ultimate or intermediate end and need not have any conception of the governing rules. It is plausible that there are, for instance, linguistic practices—practices for the use of words in a particular language—that are of this type. Learning a language from scratch, as a native speaker, does not involve setting oneself to do so. (I shall say something about this below, in section 5.) If that is right then it is not possible that the commitments incurred by participation in such practices derive from having participation as an end. We need another account of the normative dimension of practices.

The proposed explanation in any case does not adequately account for the fact that behaviour on the part of a participant that is not in accordance with a rule admits of criticism simply on account of being in breach of the rule. The criticism is of a type that might be called *agent-criticism* since it implies that the participant is in the wrong, or proceeding incorrectly, or in some way inappropriate, in relation to a standard to which it is not open to him or her to be indifferent given the activity in question. A player of football, who is in breach of a rule by trying to keep the ball in play when it has crossed the goal line, has not, simply on that account, flouted any principle of means-end rationality since no end that the player has, be it simply continuing in play or serving some ulterior end of playing, need require him to avoid trying to keep the ball in play in that situation. The criticism that is appropriate cannot, therefore, be explained in terms of a failure of means-end rationality. A natural way to describe the action would be to say that the player has made an illegitimate move or simply done something he is supposed not to do in this game. The language of

commitment gives us a way of saying what that comes to: it is a matter of being in breach of a rule and thus in breach of the commitment to follow the rules. On the view I am proposing it is simply the fact of being a participant with a sense of how one is supposed to go about things that makes sense of compliance. There might be other considerations that make sense of being a participant, but that is another matter.

Why should we suppose that the commitment to following the rules of a practice has a shape analogous to that of the commitments incurred by beliefs and intentions? To recall, the idea is that the commitment, at least a first approximation, amounts to its being the case that one ought to avoid continuing to participate in the practice while not following its rules. One reason is implicit in what I have said about the etiquette practice: the mere existence of the practice does not entail that anyone ought to follow its rules. Even a participant is not in a position such that just in virtue of being a participant he or she ought to follow the rules. But the point can be made more vivid with reference to the possibility of there being practices in which no one ought to participate. It is easy to imagine nefarious practices in which no one ought to engage. Participation in such practices still incurs commitments. If this seems odd I suspect it is because the idea of a commitment is being tied too closely to morality. The commitments incurred merely by participating in a practice are not, or need not be, moral commitments. (There is an interesting question as to the relationship between practices and morality.)

An important qualification needs to be made that explains why what I have said about the shape of the practice-grounded commitments is only a first approximation. Practices do not exist in isolation and so can be legitimately assessed from, for instance, an ethical point of view. Suppose that a person is a participant in a corrupt practice instituted by a corrupt institution. By my account such a person incurs a commitment to following the rules of the practice. That is to be understood as meaning that the person ought either to follow the rules of the practice or cease to be a participant. But surely it might be that the person ought to remain in the institution and subvert it and its rules from within. If that is so then the character of the commitment needs to be more fully specified. It amounts to its being the case that one ought to avoid continuing to be a participant and not following its rules, *in the absence of overriding reasons to do otherwise*. So far as I can see no such qualification is required in connection with the commitments incurred by beliefs and intentions.

Some practices like games are governed by rules that have been expressly designed and instituted. Some might be called informal because they are governed by rules that participants do not represent as rules and which no one has expressly instituted. Imagine a group of friends sharing an apartment.⁸ Their living together is more than just sharing a space. They often eat together, they share cleaning tasks, and they take turns to buy food and other things that are needed. We may imagine that in this group there arises a practice that could be described as informally undertaking to do something. They do not formally promise to do things. What happens is that in suitable contexts a declaration that one will, say, buy some food, counts as an undertaking to buy some food. Suitable contexts include ones in which a question arises as to who might buy food. Undertakings are in effect moves in a practice governed by a rule to this effect: if in a suitable context you say you will do something then this counts as an undertaking, so do what you undertake unless there are countervailing factors. Countervailing factors include (i) being prevented, (ii) there being an overriding reason for you do something else incompatible with doing the thing in question, and (iii) being released by the others from doing it. To say that there is this rule is to say that such a rule is at least tacitly acknowledged. In this sphere *phenomenalism* reigns: rules exist through being acknowledged and treated as rules. But they may be only tacitly acknowledged. Tacit acknowledgement is a matter of there being shared dispositions within the community. At least the following are relevant: (a) dispositions to treat the utterances in question as undertakings and therefore to expect that the person making the utterance will do the thing undertaken, absent countervailing factors, if for no other reason that that it is ‘the thing to do (the thing one is supposed to do)’; (b) dispositions on the part of those making undertakings to do the thing undertaken, absent countervailing factors; (c) dispositions to make critical evaluative judgements of failures to do the thing undertaken, despite the absence of countervailing factors or to view such criticism as appropriate; (d) dispositions to acknowledge that criticism is in order when, despite the absence of countervailing factors, one has not done something one has undertaken. These dispositions are imbued with an evaluative stance. The evaluative stance amounts to a sense that there are such-and-such right and wrong ways of behaving in the matter of making and responding to undertakings, and

⁸ This is an example I have discussed previously, in Millar 2004: ch. 3.

reacting to failures to carry out undertakings. The rule as formulated by the theorist should reflect the evaluative stance.

How should we make sense of the notion of governance by a rule in cases such as the one I have been considering? What is crucial for cases of governance by a rule, and thus for following a rule, is that participants should be to some degree sensitive to what the rule requires particular cases. They can be thus sensitive, without representing the rule to themselves, through having dispositions such as those I have described. They must have a sense of the right and wrongs associated with the activity that constitutes the practice and at least to some significant degree act accordingly. They will recognise of certain moves, which are or would in breach of the rule, that they are moves to be avoided, and they will to some significant degree aspire to avoid them.⁹

5

I suggested earlier that it is plausible that there are linguistic practices participation in which is, or need not be, viewed as an end by participants. This was important for the issue of whether it is feasible to represent the commitments incurred through participation in practices as commitments incurred through setting oneself to achieve certain ends. In this section I sketch a conception of certain linguistic practices. In the next I say something about the epistemological significance of these practices.

An attractive idea is that using words with specified meanings is engaging in a practice in the sense in play here.¹⁰ Since ‘white’ means *white* applications of the word ‘white’ to things are correct, in the sense of yielding true statements, if and only those things are white. Given that ‘white’ has those correct application conditions it is plausible that its use is governed by this rule: when using ‘white’ respect those conditions of correct application. This does not require that one should apply the term ‘white’ to things if and only if they are white. If knowing that I speak falsely I tell my elderly aunt that the walls of her great nephew’s sitting

⁹ Ryle says of someone who knows how to play chess, ‘His knowledge *how* [to play chess] is exercised primarily in the moves that he makes, or concedes, and in the moves that he avoids or vetoes’ Ryle 1963: 41. He adds, ‘So long as he can observe the rules, we do not care if he can also formulate them’. Compare also Wittgenstein’s remarks leading up to the claim about obeying a rule blindly (1958: Part I, §§ 206-219). Among those who have tried to make sense of rule-following in the absence of any representation of the rules is Sellars 1963 and Bennett 1964.

¹⁰ The conception to be introduced here is set out in Millar 2004: ch. 6. An earlier version is Millar 2002.

room are white believing that she would like this to be so then I lie. What I say is incorrect in the sense that I speak falsely, but there is another sense in which my use of the term ‘white’ is perfectly correct for it is entirely in keeping with the term’s meaning. I wanted to tell my aunt that the walls of her great nephew’s sitting room are white and used just the right word for that purpose. I respect the word’s conditions of correct (= true) application since I meant to say that the walls of my sitting room are white, needed to use a word with conditions of application such that it would correctly apply to something if and only if that thing is white, and did use just such a word.

Similar considerations apply in cases in which someone speaks falsely but aims to speak the truth. A doctor might say, sincerely, that a patient has flu when the patient does not have flu but has meningitis. This application of the term ‘flu’ is incorrect in the sense of being false, yet the doctor is using the term correctly in the sense that his use is in keeping with the term’s meaning. It is the right word to use to say in English that a person has flu. The doctor’s statement respects the conditions of application of the term in that he meant to say that the patient has flu, needed to use a word that correctly applies to somebody if and only if that person has flu, and did use just such a word.

What would be a case of not respecting the conditions of application of a term? The patient in Burge’s famous example is a case in point.¹¹ He applies the term ‘arthritis’ to any painful condition of the limbs, when in fact it applies to painful conditions due to inflammation of joints. He uses the word as if it correctly applies to a condition if and only if that condition is a painful condition of the limbs. So he uses the term as if its conditions of application are other than they are. That is to fail to respect the actual conditions of its correct application.

Notice that, as the example of lying shows, applications of a term that are incorrect in the sense of being false applications might be correct in the sense of being in keeping with the conditions of correct (= true) application. Likewise, applications that are incorrect in the sense of not being in keeping with the conditions of correct (= true) application might be correct in the sense of being true applications. Burge’s patient can illustrate the latter case. He uses the term as if the conditions of its correct (= true) application are other than they in fact are. Yet, despite this, he may on occasion correctly (= truly) say of someone, ‘He has arthritis’.

¹¹ Burge 1979.

Granted that there are rules for terms of the sort under consideration—rules that enjoin us to use terms in ways that respect their conditions of correct application—it does not seem plausible that individuals who use terms represent the rules for using those terms to themselves. (Indeed, even theorists have not in general represented rules for the use of terms as I have done.) But we can, I think, make sense of the idea that users of terms tacitly acknowledge such rules. This would be manifested by their being disposed to treat breaches of rules as mistakes, and by their willingness to adjust their usage of a term on learning that it is not in keeping with its conditions of correct application.

In view of Burge's example we had better accommodate the possibility that people can be participants in a practice for use of a term while systematically using that term incorrectly (= out of keeping with its conditions of correct application). Burge's patient is in the business of using the term 'arthritis' in keeping with its conditions of correct application. He aspires to use it correctly, as opposed to using it in keeping with some stipulated meaning that he has adopted. That is why he is willing to adjust his use on learning, in effect, what the conditions of its correct application are. His problem was that until this learning occurred he did not know what the conditions of correct application are. Having this problem he was not in a position to recognise that certain applications of the term 'arthritis' are incorrect (= false) or are misuses in the sense of uses that are not in keeping with the terms meaning.

Because of cases like that of Burge's patient, we need to accommodate the idea that someone can be a participant in a practice yet have dispositions to use the term that are at odds with the rule governing its use. The result will be that the extension of the notion of a participant will be indeterminate. There will be cases that clearly fall within the extension and cases that are close to doing so but where there is no fact of the matter as to whether they do or do not. Burge's patient counts as a participant in something like the way in which a novice footballer counts as a player prior to having mastered the rules. Both aspire to do the done thing, and are some way along the road to having a grasp of what it takes to do that, but they get some things wrong.

It would be wrong to say that all participants in practices of using terms aspire to use those terms correctly (= in keeping with their meanings) or that they view participation in practices for the use of words as an end. This would require a degree of reflectiveness that very young participants do not have. Nonetheless, very young children taking their first steps in language count as participants simply because they are being brought up into it and making some headway in respect of what the practice requires. This has a bearing on whether the normativity of participating in practices can be explained in terms of instrumental rationality.

Very young children do not view their use of a language either as an end in itself or as a means to the end of communication. Indeed, it is not so very bold to suppose that most of us most of time do not view our language as a means to the end of communication any more than we view breathing as a means to living. We do it anyway. Our participation in linguistic practices is a given that need have nothing to do with a desire or intention to participate or with any ulterior aim of participating. It is something into which one is initiated irrespective of any choices one makes and aims one has.

Linking the preceding considerations with the general account of practices sketched earlier, we have it that participants in practices for the use of terms incur a commitment to using those terms in keeping with the rules for employment. Among those rules are rules to the effect that one should respect the conditions of correct application of this or that term.

6

Linguistic practices nicely illustrate the epistemological significance of practices. Someone I have never met—an assistant in a gallery selling art works—utters in fluent English the words, ‘This one costs £100’, pointing at a ceramic dish in the gallery. I know nothing of this person beyond what I can gather from the present situation. Yet I recognize that he is telling me that the dish picked out costs £100. To explain this we need to explain a number of things including the following: (A) how it can be that uttering those words in such a context indicates that one is saying of something picked out that it costs £100; (B) how it can be that saying this sort of thing of something picked out, in the manner it is said here and in a context of this sort, indicates with a high degree of reliability that one means to say of the thing picked out that it costs £100; (C) how it can be that saying, and meaning to say, of something picked out that it costs £100 can be an act of telling those addressed that the thing picked out costs £100.

These explananda are not, in themselves, epistemological but they figure in a plausible practice-theoretic account of the epistemology of understanding utterances. To explain (A) on this account we need the idea that the meanings of words are fixed by practices for using those words. Practices for the uses of terms are governed by rules prescribing respect for their conditions of correct application. What people say, if anything, by a use of words, is fixed by what those words in combination mean. To explain (B) we need the idea that just because there are practices for the use of words it will be the case that, absent rare countervailing factors, when people say something intelligible and to the point and in a manner that

manifests linguistic competence, they mean to say it.¹² That is because thanks to the relevant practices, absent rare countervailing factors, such people use words in keeping with the rules for their use. In the case described the assistant's saying of the dish that it costs £100 can clinch it that he meant to say that the dish costs £100. Clearly, not all acts of saying are acts of telling. Expressing an opinion, for instance, is distinct from telling, since in telling one gives it to be understood that one knows the thing one tells and intends that the addressee should come to know it from the telling. To explain (C) we need the idea that there is a practice of informing by telling such that saying something in a certain manner, under certain conditions, counts as telling that thing to those addressed.

The explicitly epistemological part of the story includes an explanation of the following: (D) how I can tell that by uttering the words he does the assistant is saying of the dish picked out that it costs £100; (E) how I can tell that by saying this the assistant meant to say it; (F) how I can tell that in saying this, meaning to do so, the assistant was telling me that the dish costs £100.

The quick and rough answer to all of these matters is that initiation into the relevant practices, including those described above, results in the acquisition of certain complex recognitional abilities that are exercised in telling of utterances that they are acts of saying to us, and indeed of telling us, that the thing picked out costs £100. Exercise of such abilities results in knowing that one is being told concerning the thing picked out that it costs £100. In the example under consideration, I know that I am being told that the dish costs from the words uttered but I do not take in these words simply as words but hearing their utterance as an act of saying what they say. I tell what I am being told from what is being said to me. In this process the crucial point for present purposes is that within the practices in play what one is supposed to be doing by saying such a thing in such a manner and in such a context is telling me that the dish costs £100. I am not suggesting that I engage in reasoning (i) that he is saying to me that the dish costs £100 in such and such a manner and context; (ii) that by so saying he is supposed to be telling me that the dish costs £100; (iii) not easily would it be that he is supposed to be telling me that that the dish costs £100 and yet he is not telling me this; (iv) so he is telling me this. My acquisition of knowledge of what he is telling me is phenomenologically immediate. Yet it is not blind for it is informed by an understanding of

¹² That is not say that they believe the things they say. People can speak out of bitterness, saying things they regret and do not really believe. That is consistent with their meaning to say what they did when they said it.

what one is supposed to be doing when saying such a things in this manner and in such a context. If someone were to raise a question as to whether the assistant did tell me that the dish cost £100 I would rightly wonder what else he could have been doing because this is what anyone uttering those words in the way he did in the relevant context is supposed to have been doing. This understanding does not require that one hold some complex theory about language and meaning. It does require that one should have some sense of what one is supposed to be doing by using certain words in certain ways and is attuned to the significance of what one is supposed to be doing by so using those words, that is to say, to its significance with respect to what the speaker is actually doing.

Why should we suppose that considerations about what people are supposed to be doing by using words in the way they enter the picture at all? Because this is a plausible, and I suspect the most plausible, explanation for the grounds of our judgements as to what people are doing with words.

It seems especially clear that in the case of linguistic practices we do not work with empirically grounded generalizations about how words are in fact used. Of course, we acquire habits of expectation via the impacts of myriad uses of words upon us. And in grasping what a word means we presume that by-and-large people will use it in keeping with its meaning. But that presumption is one we are committed to accepting in regarding people as speakers of a particular language, rather than a generalization acceptance of which is justified by evidence. It is an instance of a type of generalization to which we are committed by viewing people as engaged in certain types of activity—the type that the theorist identifies as being practices. There are indicators that people are so engaged and accordingly will by and large do what they are supposed to when so engaged. From this perspective we should not think of the generalizations in question as ones for which we need independent evidential support. They are rather part of the understanding that we work with in our interactions with those who participate in the relevant practices.

An alternative to my view is the idea that our attributions of acts saying and telling to others is grounded in judgements as to what we would be doing if we used words as they are. I cannot myself see why judgements about what we would be doing should be taken to be indicative of what others are doing unless we are all participants in practices for using the relevant words and making the relevant speech acts. And if we are then we don't need a detour through judgements about what we would be doing.

References

- Bennett, J. (1964). *Rationality*. (London: Routledge and Kegan Paul).
- Burge, T. (1979). 'Individualism and the Mental', in P. A. French, T. E. Uehling, Jr., and Howard K. Wettstein (eds.) *Midwest Studies in Philosophy, iv*. (Minneapolis: University of Minnesota Press: 73-121.
- Foot, P. (1972/1978). 'Morality as a System of Hypothetical Imperatives', *The Philosophical Review*, 81: 315-16. Reprinted in Foot's collection *Virtues and Vices* (Oxford: Blackwell, 1978). Page references are to this reprint.
- Millar, A. (2002). 'The Normativity of Meaning' in A. O'Hear (ed.) *Logic, Thought and Language* (Cambridge: Cambridge University Press).
- Millar, A. (2004). *Understanding People: Normativity and Rationalizing Explanation* (Oxford: Clarendon Press).
- Millar, A. (2010). 'Knowing From Being Told' in Haddock, Millar and Pritchard (eds.) *Social Epistemology* (Oxford: Clarendon Press).
- Pritchard, D., Millar, A., and Haddock, A. (2010) *The Nature and Value of Knowledge* (Oxford: Clarendon Press).
- Ryle, G. (1963) *The Concept of Mind* (Harmondsworth, Middlesex: Penguin Books).
- Sellars, W. (1963). 'Some Reflections on Language Games', in Sellars's *Science, Perception and Reality* (London: Routledge and Kegan Paul).
- Wittgenstein, L. (1958). *Philosophical Investigations* Second Edition. (Oxford: Blackwell).