The Epistemological Significance of Practices

Alan Millar

1. The Problem

Much of our knowledge concerns the thoughts and actions of others. We often know what people are doing and why they are doing it in the light of propositional attitudes we know that they have. On the basis of known attitudes we often make reasonable predictions about what people will think or do in this or that circumstance. Having arranged to meet a friend at a particular time and place it is reasonable for you to expect this friend to be at that place at that time. It is part of the arrangement that each party to it takes the other to intend to meet at the agreed time and place and on that basis expects the meeting to take place. Such understandings and predictions of behaviour are routine because we routinely know about what people believe, desire, or intend, or how they feel about this or that type of circumstance. Often they simply tell us about such matters. Yet there are countless occasions when we find people's thought or action intelligible, or anticipate what they will think or do, or are at least unsurprised by what they think or do, despite our having little if any information about their attitudes other than what we can gather from their situation and non-verbal behaviour. Watching a football match I see a player in one team trying to avoid being too close to two players in the opposing team. I 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 having no specific information about any of the propositional attitudes of the players apart from what I may gather from their circumstances and their non-verbal behaviour. Other examples easily reinforce the point. Suppose I receive an invitation to give a talk at some academic institution. All I receive is a brief email, from someone I do not know, telling me about the nature of an event that is being arranged and asking me to give a presentation. I reply saying that I would be happy to come. I readily understand that the email is an invitation. I expect that the sender will ensure that the event takes place as described. I also expect him or her to take from my reply that I will come as arranged. My response is reasonable though made in the absence of evidence directly concerning the personal traits or propositional attitudes of my interlocutor beyond what I can gather from the email. What is puzzling about such situ-

ations is how it can be that despite the paucity of information we have about people's attitudes—information that is specifically about them—we still seem able to arrive at correct or at least reasonable understandings and predictions of much of what they think and do. These understandings and predictions, it seems, have to be in the light of their attitudes yet we seem to know so little specifically about their attitudes.

One response is to deny that there is anything puzzling here. After all, it is not to be expected that whenever the behaviour of people makes sense to us we have information deriving from them, or anyone else, specifically about their attitudes. We are not surprised to see lots of people on a rail station platform at rush hour. Though we know nothing about them as individuals, it makes sense that they should be there because it is to be expected at a time when lots of people will be returning home from work. A strand of empiricist philosophical tradition encourages us to think that we employ a theory—a body of rough generalizations representing plain matter-of-fact regularities in human behaviour—which we are justified in accepting and which we apply to particular individuals and circumstances.[1] On this line of thought, in the scenario of the rail station we might rely on a generalization to the effect that during rush hours there will be lots of people on rail station platforms waiting for trains to take them home. In the scenario of the invitation I might exploit, 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. The trouble is that it is far from obvious that in all cases we have suitable generalizations available. Someone utters a sentence in my presence. I understand what the person is up to because, among other things, I identify the person's utterance as one of *telling* me that something is so, as opposed, say, to voicing an opinion or giving me advice. What generalizations are at my disposal in a case like this? I respond to a host of cues that I would be hard put to specify except in the most general terms. No doubt the content and context of the utterance and the manner of speaking are relevant, along with much else besides. But how am I to spell out, for instance, the determinate form of the manner of speaking that is relevant? 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, or as a case of telling whatever it does tell. We respond to a total *Gestalt* just as we do when recognizing a face as the face of some individual. Our response is recognitional rather than based on assumptions about features of the utterance. But if at the level or belief or judgement we do not take in which features are

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

relevant to its being the speech act that it is then we can hardly apply a covering generalization pertaining to those features.[2]

Even when generalizations are available there is a question as to whether we are justified in thinking them true. In relation to the invitation scenario, it is at least open to doubt that experience, including that of the testimony of colleagues, will have supplied me with evidence sufficient to justify me in thinking that when people at academic institutions issue invitations they are likely to ensure that the event in question will take place unless unforeseen circumstances prevent this. I might never even have entertained such a thought. It seems more likely that I would have acquired a disposition such that on being given an invitation I would expect the event in question to take place, in the absence of reasons to the contrary. We passively acquire such dispositions. They are undoubtedly shaped by our experience, but that is not say that experience supplies us with a reason to accept a corresponding generalization—a reason for which we accept such a generalization.

The view to which I have been reacting—a version of the theory-theory of understanding others—is usually contrasted with simulation theory, so one might wonder whether the latter would fair better in relation to the problem. The core idea of simulation theory is that we make sense of what people think and do on the basis of simulations undertaken as if in their shoes.[3] We ascribe to the person attitudes and actions reflecting how we would think and act in their situation, allowing for relevant differences between us. I shall return to simulation theory later. For present purposes the point that matters is that even if some version of simulation theory were correct it would not by itself solve the problem. We would still need, for the cases under consideration, an explanation of why a simulation should be expected to succeed in mimicking the relevant pattern of attitudes and actions of those we seek to understand and of how we come by the resources that enable us to put ourselves in their place.

An important factor in understanding and predicting the behaviour of others is that we grasp what people have reason to think or do, or are committed to thinking or doing, in the circumstances. To take a simple example: it makes sense that a person crossing a road should quicken her step if a car is rapidly approaching because it makes sense, given her orientation, that she should have seen the car and recognize that she had better get out of its way. In the present discussion I shall focus on a special type of case in which our understandings

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

3 For useful collections of articles, see Davies and Stone 1995a and 1995b. For a developed defence of simulation theory, see Goldman 2006.

and expectations depend on there being certain actions that are, so to speak, *the done thing*—the thing one is supposed to do—because of the character of the activity in which the person is engaged. A natural way to fill this out invokes the notion of a practice in the sense of an essentially rule-governed activity. In pursuing this theme I shall build on and develop lines of thought in *Understanding People* (Millar 2004).[4]

2. Towards a Solution

Practices have both explanatory and epistemological significance. A game of football is a practice in the relevant sense. The behaviour of players is shaped by the aim of the game—to win by scoring more goals than the opposing team in accordance with the rules of the game. Much of what players do in the game is in conformity with its rules because they have a more or less adequate grasp of what the rules demand and, being players, comply to a significant extent with what the rules demand. So the nature of the practice—in this case playing the game—has a substantial role in explaining why players act as they do while playing. Both players and observers of the game make sense of what is going on and can form reasonable expectations about what is going to happen at various stages of the game even in the absence of knowledge of the views, styles and techniques of particular players. In all of this they, of course, presuppose that the players have certain propositional attitudes, but their epistemic route to the attitudes is not independent of an understanding of what the rules require.

In the case of football rules are instituted and represented verbally so that they can function as instructions. Not all practices are governed by rules that are represented in language and acknowledged explicitly as rules. Imagine a group of friends sharing an apartment.[5] 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 surrounding what could be described as *informally undertaking to do something*. What happens is that in suitable contexts a declaration that one will, say, buy some food, counts as an undertaking to

4 Given what I have said so far it perhaps goes without saying that I owe a very large debt to the work of Donald Davidson. For some relevant discussion of Davidson, see Millar 2004: ch. 1. The idea that there are essentially rule-governed activities and that they are important for understanding human conduct is not new. See, for instance, John Searle's discussion of what he calls institutions (Searle 1969). My own treatment is distinctive in according a central role to commitments as set out in Millar 2004 and below in section 3.

5 This expands further on an example I have discussed previously, in Millar 2004: ch. 3.

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 rules to this effect: (1) if in a suitable context you say you will do something then this counts as an undertaking, that is, something you give others to understand that you will do, unless there are countervailing factors; (2) do that which you undertake, unless there are countervailing factors. Countervailing factors include (i) being prevented, (ii) there being an overriding reason for you to do something else incompatible with doing the thing in question, and (iii) being released by the others from doing it.[6] In this sphere a kind of *anti-realism* holds: rules exist only through being acknowledged as rules and complied with to some significant degree. The less acknowledgement there is across the community the less clear it is that the rules exist, but the acknowledgement might be only tacit. Tacit acknowledgement of rules is a matter of acknowledging of certain actions prescribed by the rules that they are *to be done*, acknowledging of others that are prohibited by the rules that they are *not to be done*, and being willing by-and-large to do, and approve of, the former and avoid, and disapprove of, the latter. In place of recognition of rules in full generality there is an ability to recognize and appropriately respond to what is and what is not to be done in particular cases. Those in the imagined community will be willing to treat the utterances in question as undertakings and, accordingly, will expect those making such utterances to do what they undertake to do, absent countervailing factors, if for no other reason than that it is 'the thing to do (the thing one is supposed to do)'. They will be prone to judge adversely failures to do the thing undertaken, when there are no countervailing factors. They will be prone to acknowledge that criticism is in order when, despite the absence of countervailing factors, members of the community have not done something they undertook to do. When they make undertakings themselves they will be willing to do the thing undertaken, absent countervailing factors, just because they have undertaken to do it. What all this amounts to is that governance by a rule—following a rule—can be a matter of being disposed to be appropriately responsive to what the rule requires in particular cases. One can be thus sensitive, without representing the rule to oneself, through being prone to the sorts of reactions I have described.[7]

6 That a practice such as this should come into being might well have to do with the existence in the wider society of a more formal practice surrounding promising and of a morality encouraging respect for others. Even so, it is reasonable to think of the practice in the apartment as their practice since its rules have a grip on them against a background of their habits of behavior and their willingness to act cooperatively in this particular setting.

7 Gibert 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

Other, less complex, examples abound. There is a British practice that surrounds queuing for a bus governed by a tacitly acknowledged rule to the effect that you do not get on the bus until others who were at the bus stop before you have done so.

If invoking the notion of a practice is to contribute to addressing the problem posed by paucity of information specifically about individuals we shall need to be assured that we can make clear sense of the idea that there are practices as understood here. That requires an explanation of what my talk of the normative dimension of practices amounts to. I turn to this in the next section.

3. The Normative Dimension of Practices

I find in discussion that the idea that there are practices that have an intrinsic normative dimension is viewed by some as being rather obvious and by others as being highly controversial. Where it is accepted there can be disagreement about how the normative dimension should be conceived. Here I provide an outline of an account and address sources of scepticism about the very idea of a practice.

Both the explanatory and the epistemological roles of practices depend on how participants in a practice relate to the rules of the practice. The central claim in the proposed account is that participants have a certain normative connection to the rules of the practice in that just in virtue of being participants they incur a (normative) commitment to following—being governed by—the rules of the practice. This gives to practices a normative dimension that is essential to what they are.

The sense of 'commitment' here is not the psychological sense which has to do with being determined or resolved to do something. It is a normative sense that is modelled on a conception of the commitments incurred by beliefs and intentions.[8] 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.

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'. Among those who have tried to make sense of rule-following in the absence of any representation of the rules are Wilfred Sellars 1963 and Jonathan Bennett 1964.

8 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.

(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, but that cannot be exactly right. By (3) the commitment, so understood, would amount 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 clearly too strong since 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 commitment we are looking for must be to believing that q *if one gives any 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.

There is a sense in which withholding judgement on whether q is also, in the relevant sense, 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 I withhold judgement on something the truth of which is part of the picture that I have.[9] Note that there are two ways to *discharge* the commitment: by believing that q or by giving up one of the beliefs that jointly incur the commitment.[10] I might not be in breach of the commitment if, having believed that p, and that if p then q, I do not believe

9 One might think that the commitment to which I have 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. 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. For further discussion of how these commitments relate to rationality, see Millar 2004: ch. 3, section 2. My conception of the (wide scope) shape of the commitments under consideration has been much influenced by John Broome's work (notably, Broome 2002).

10 One can avoid being in breach without discharging by never giving any verdict on whether q.

that q, since I might avoid a breach by giving up believing that p or giving up believing that if p then q.

Commitments incurred by intentions—specifically, commitments to taking the means necessary to carry out our intentions—have a closely analogous structure. If I intend to ϕ, and if ψing is necessary if I am to ϕ, then I incur a commitment to ψing. So if I have the intention I ought to avoid being such that I retain that intention yet never get around to ψing. Here too there are two ways to discharge the commitment, by ψing or by giving up the intention to ϕ. (In this case there is no way to avoid being in breach short of discharging.)

The commitments incurred by participation in practices have an analogous structure. Participating in a practice, G, incurs a commitment to following the rules of G. To a first approximation that amounts to it being true of each participant in G that he or she ought to avoid continuing to be a participant while not following the rules of G. (But see below for an important qualification.) As before there are two ways to discharge the commitment: (i) by ceasing to participate in the practice and (ii) by complying with the rules.

In the case of the commitments incurred by beliefs and intentions it is clear that there is reason not to represent the commitments as absolute obligations do the thing to which one is committed: sometimes the right thing to do will be to alter the commitment-incurring condition; one will not be in breach of any obligation simply in virtue of not giving any verdict on something entailed by what one believes. 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?

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).

I think there is something right about this. As such it can be turned into a problem for the theory of practices since etiquette can plausibly be taken to be a practice governed by a cluster of rules for behaviour while dining, and

so forth.[11] 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 practice of etiquette gives one a reason to follow its rules. Incurring a commitment to following the rules favours or recommends following the rules. Yet, even if this is right, the key point that Foot is making stands: to anyone considering if it matters whether one behaves in keeping with rules of etiquette the mere fact that there are these rules does not show that one ought to comply *even if one has up to this point participated in the practice*. 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 on my account we can accommodate Foot's key point. The normative dimension of a practice 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.

Even if these points are conceded one might persist in feeling that there is a straightforward sense in which, for instance, if one is playing in a game of football one ought to follow the rules. That seems right but, nonetheless, the *ought* in such a thought is weaker than that in play in my account of commitments. The latter I take to be the deliberative *ought* that expresses the outcome of deliberation over what to do *all things considered*. The weaker *ought* is apt for expressing something that there is some reason to do, but which might not be, all things considered, the right thing to do.

Another sceptical line of thought that might be advanced at this point has it that the appearance of there being a normative dimension to practices can be explained in terms of instrumental rationality as follows.

> To understand playing games that are defined by rules we just need

11 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 instrumental rationality. A player of a game of football intends to play the game and so is committed to doing 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 latter are grounded in intentions to participate, so there is no need to posit a special kind of normativity that is specific to practices. The rules themselves simply set out modes of behaviour such that if players do not accord with those modes of behaviour they incur a sanction. So while there is a sense in which players in a game incur commitments to comply with the rules, this commitment is explained in terms of general considerations about instrumental rationality; it does not require us to assume that practices have an intrinsic normative dimension.

The suggested account might seem 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 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. Linguistic practices, including practices for the use of words in a particular language are of this type. Learning a language from scratch, as a native speaker, does not involve setting oneself to do so. If that is right then it is not plausible that the commitments incurred by participation in such practices derive from having an intention to participate. (I shall say more about this in the next section.)

Might a more nuanced instrumentalist response do the trick? Surely participants in linguistic practices have a reason to comply with the rules, irrespective of whether they have participation as an end. This reason might be viewed as grounded in a desire or need to communicate with others. The question though is not whether there is such a reason—let as grant that there is—but whether it can account for the appearance of there being an intrinsic normative dimension to practices. The proposed explanation 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 a type of *agent-criticism* since it implies that the participant is in the wrong, or proceeding incorrectly, or in some way inappropriately, in relation to a standard to which it is not open to him or her to be indifferent given the

activity in question. Where practices are concerned the standard is provided by compliance with the rules. 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, or reason for him or her to play, 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 with the rules. There might be other considerations that make sense of being a participant, but that is another matter.

A qualification is needed because it is easy to imagine nefarious practices in which no one ought to engage. 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, in the strong deliberative sense explained earlier, either to follow the rules of the practice or cease to be a participant. But surely it might be that the person, in this sense, 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 while not following its rules, *in the absence of overriding reasons to do otherwise*. No such qualification is required in connection with the commitments incurred by beliefs and intentions.

4. Practices of Using Words

What I have said about the commitment incurred by participating in a practice will, I hope, make it plausible that the notion of a commitment is not highly theoretical one of which only theorists have some grasp. Players in a game of football have a sense of how they are supposed to behave in virtue of being players. In the example of sharing the apartment, even in the absence of explicitly formulated rules, the householders have a sense of how someone who

undertakes to do something is supposed to behave and a sense of appropriate responses to failures to keep undertakings.

In this section I shall focus on a type of linguistic practice—practices of using words in a natural language in keeping with particular meanings.[12] These are of particular interest in the present context because the problem posed by paucity of information specifically about particular individuals is at its most acute. We can understand strangers by grasping what they say and the point of their saying it. It is not at all plausible that we rely on generalizations representing brute regularities in the use of words. Even so, a practice-theoretic account will look plausible only if we can make sense of the idea that there are linguistic misuses of words—uses of words that are not in keeping with rules for their use. Here in briefest outline is the direction in which I think we should go.

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. So we may think of the use of the word to mean *white* as a practice—an activity that is essentially governed by this rule. It is of some importance that this view 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 my sitting 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 is perfectly correct for it is entirely in keeping with the term's meaning. I wanted to tell my aunt that the walls of my sitting room are white and used just the right word for that purpose. I respect its 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.[13]

12 The conception to be introduced draws upon Millar 2004: ch. 6. An earlier version is Millar 2002.

13 While denying that word meaning is intrinsically normative, Paul Horwich (1998: ch. 8) takes it that language has a normative aspect expressible in terms of such claims as that if 'dog' means dog then it ought to be applied to, and only to, dogs. He thinks that 'ought' here is explained in terms of having truth as an end. The idea that the normativity associated with meaning is to be understood in terms of such ought-claims derives from Kripke 1982, where the normative aspect is taken to be essential to meaning. His view of the shape of the normative dimension to meaning has been enormously influential on subsequent discussion. Attacks on the idea that normativity is essential to meaning, when it is understood in this way, are taken to undermine any view that takes meaning to be essentially normative, as in Hattiangadi 2006. For a useful survey, see Glüer and Wikforss 2009.

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 not thereby misusing the term. He uses it 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 correct 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 Tyler Burge's (1979) famous example is a case in point. 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. In a weak sense this patient respects the conditions of true application of the term; after all, he takes it to stand for the condition arthritis and he aspires to use it in keeping with that meaning. But in another, stronger, sense he fails to respect the conditions of true application of the term because he has a false conception of what arthritis is. 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. For another example, imagine a child who has learned what foxes look like from illustrations in books but has never seen a dog. The first time he sees a German Shepherd dog he calls it a fox. Here too we have a case of misusing, not merely misapplying, a term. The term means *fox* but it is being used as if it meant something like *fox-like animal*, where likeness is conceived loosely. That counts as a failure to respect its conditions of 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 illustrates the latter case. He uses the term in a manner that is not in keeping with its conditions of correct (= true) application. This flows from a misunderstanding of the condition to which the term applies. Yet, despite this, he may on occasion correctly (= truly) say of someone, 'He has arthritis'.

Linguistic practices of the sort under consideration serve to highlight a theme on which I touched in section 2—that some rules are only tacitly acknowledged. Granted that there are 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, 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, for instance, 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.

Could we not account for use simply in terms of the dispositions of subjects to use these terms? The trouble with the latter suggestion is that it fails to accommodate the possibility that people can systematically misuse a word while using it in such a way that it counts as meaning what it does. Burge's patient is in the business of using the term 'arthritis' in keeping with its conditions of correct application in that 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 when he arrives at a better understanding of its conditions of correct application. Prior to this his dispositions to use the term were liable to issue in incorrect (= false) applications yet he still counted as using it such that it means *arthritis*. Were those speaking English to render what he says when he says things like, 'He has arthritis' they would be right to do so using the word 'arthritis', without some qualification like 'in his peculiar sense'. I think we can best make sense of this phenomenon if we see the meaning of the term, even as he uses it, as being tied to a practice that is independent of him. He counts as using the term in such a way that it means *arthritis* because he is a participant in the practice governed by the rule of using it in keeping with its conditions of correct application. He counts as a participant in the relevant practice in something like the way in which a novice footballer counts as a player prior to having mastered the rules. Both, as it were, 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 both get some things wrong.

In the next section I highlight some epistemological implications of the preceding conception of linguistic practices.

5. Epistemological Ramifications

As I remarked above, linguistic practices form a class that is of particular interest given the problem identified at the beginning of the present discussion. For we routinely make sense of the utterances of people about which we have

no specific information beyond what we can take from their present utterances, the context of those utterances, and their non-verbal behaviour. Yet understanding their utterances is inextricably bound up with finding intelligible what they are doing with the words they use and that involves viewing the utterances as having a background of propositional attitudes, including intentions and beliefs.

In thinking about words we think about what they mean and that, on the present picture, carries implications for how one who uses them is supposed to use them—is (normatively) committed to using them. Of course, we would lose any sense of how words are supposed to be used if our experience showed that they were not regularly used as they are supposed to be. But that is not to say that we represent regularities of use at the level of belief or judgement and think in terms of those representations.

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 by his use of these words he is telling me that the dish picked out costs £100. A necessary condition for me being able to tell from his utterance that he is saying that the dish indicated costs £100 is that uttering such words as this person did indicates with very high degree of reliability that the person is saying of what is indicated that it costs £100. (Compare: a necessary condition of me able to tell of this structure that it is a barn from its visual appearance is that possession of such an appearance is a very highly reliable indicator of being a barn.) What makes it the case that the reliability condition is met is that, amongst other things, there is exists a practice of using 'costs' in keeping with a rule prescribing that one respects the condition that it is correctly applied to an object and a sum of money, in that order, if and only if the object costs that sum of money.[14] However, to comprehend the utterance I need not have, and would not have, satisfied myself that the reliability condition holds. I simply recognize what is being said from the sound of the words uttered. The recognitional abilities that I exercise in this context have been acquired through initiation into the relevant practices, including that for using 'costs' in keeping with the specified conditions of correct application.

My acquisition of knowledge of what is being said to me is phenomenologically immediate; it is recognitional just as my perceptual knowledge that

14 As conceived here the reliability condition can be met even if the utterance of these words were unique, since it is a modal condition to the effect that not easily would someone utter those words in such a context and not be saying that the thing picked out costs £100.

a thing I am looking at is a barn can be recognitional.[15] Yet recognition in mature speakers is not blind for it is informed by an understanding of what one is supposed to be doing when saying such a thing in this manner and in such a context. This understanding would be manifested if someone were to raise a question as to whether the assistant did tell me that the dish cost £100. For in that circumstance I would rightly wonder what else he could have been doing given that this is what any speaker of English 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 words mean and, along with that, a sense of how they are supposed to be used.

6. Back to Simulation Theory—Briefly

Earlier I commented that even if some version of simulation theory were correct we would still need an explanation of why a simulation should be expected to succeed in mimicking the pattern of attitudes and actions exhibited by the person we seek to understand or whose thought or action we seek to anticipate and an explanation of how we acquire the resources to put ourselves in the place of that person. I shall not attempt a fully general answer to this question. I shall focus on cases in which practices are in play.

Consider again the situation in which someone says to me, 'This one costs £100'. According to simulation theory I understand this in the light of what I would be doing if I were to use these words in a similar context. It is pertinent to ask why it should be expected that if I were to tell someone, 'This one costs £100' what I would tell concerning me would be the same as what my interlocutor is telling me concerning him, when he says to me, 'This one costs £100'. The practice-theoretic account delivers an answer: my interlocutor and I are participants in practices for using the words in question. So, among other things, we use the word 'costs' in similar ways and expect other speakers of English to use them in much the way we do, allowing for divergences in understandings of the property the words pick out. So there is no problem about how I can put myself in the place of the other, since I have a practical grasp of practices in which we both participate. From this standpoint simulation theory does not displace a practice-theoretic treatment of understanding

15 Recognition and recognitional abilities are given a prominent epistemological role in Millar 2008, 2009, 2010, and in my contribution to Pritchard, Millar, and Haddock 2010.

activities of the sort conceived here as essentially rule-governed. Arguably, it requires supplementation by such a treatment.

More controversially, given a practice-theoretic account it becomes a serious question whether simulation is needed to work out what my interlocutor is up to when he says what he does. For, having been initiated into the relevant practices, I am equipped with the abilities that enable me to recognize what people are saying and doing when they utter, 'This one costs £100' and, accordingly can tell straight-off that my interlocutor is telling me that the dish picked out costs £100. My ability to do so is informed by my understanding of the business of using words in the way the speaker does. That implicates an understanding that using 'costs' commits you to using it in a manner that respects its conditions of correct (= true) application. It is not clear why it should be thought that I would need to run through what *I* would be doing if I uttered the relevant words. What I would do is in all likelihood what my interlocutor would do but that does not mean that my access to the latter is via the former.

A similar scepticism seems called for in less complicated cases. Suppose that part of the Dean's job is to endorse recommendations of appointing committees, in the absence of any reason to question such recommendations. Then if an appointing committee recommends appointment of a candidate who is clearly way ahead of other candidates, and due process has been followed, I shall expect the Dean to endorse this recommendation. The basis of my expectation is that in virtue of his office the Dean is committed to following the rules that apply to those holding this office, which in this case involves a commitment to endorsing the recommendation. Of course, if he does endorse it he would be doing exactly what I would do if I were Dean. But it is not clear why it should be thought that I would run through what I would do in order to make out what the Dean will do. I have to know what Deans are supposed to do on the matter in hand. Much of the present discussion has been aimed at shedding light on what that sort of consideration amounts to.[16]

References

Bennett, J. (1964). *Rationality*. (London: Routledge and Kegan Paul).
Broome, J. (2002). 'Practical Reasoning' in J. L. Bermúdez and A. Millar (eds.) *Reason*

16 Some lines of thought in this article have been provoked by David Papineau's scepticism on various occasions. I have also benefited from audience comments at conferences, workshops or seminars in London, Barcelona, Geneva, Edinburgh and Stirling. I am especially grateful to my colleagues Colin Johnston and Simon Hope whose comments prompted amendments to earlier versions of the text.

and Nature (Oxford: Clarendon Press, 2002), 85-111.
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.
Davies, M. and Stone, T. (eds.) (1995*a*). *Mental Simulation* (Oxford: Blackwell).
Davies, M. and Stone, T. (eds.) (1995*b*). *Folk Psychology: The Theory of Mind Debate* (Oxford: Blackwell).
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, 57-73.
Glüer, K. and Wikforss, A. (2009). 'The Normativity of Meaning and Content' in *The Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/entries/meaning-normativity/. Accessed on 2 August 2011.
Goldman, A. I. (2006). *Simulating Minds: the philosophy, psychology, and neuroscience of mindreading* (Oxford: Oxford University Press).
Hattiangadi, A. (2006). 'Is Meaning Normative?' *Mind and Language*, 21, 220-40.
Horwich, P. (1998). *Meaning* (Oxford: Clarendon Press).
Kripke, S. (1982). *Wittgenstein on Rules and Private Language* (Oxford: Blackwell).
Millar, A. (2002). 'The Normativity of Meaning' in A. O' Hear (ed.) *Logic, Thought and Language* (Cambridge: Cambridge University Press), 57-73.
Millar, A. (2004). *Understanding People: Normativity and Rationalizing Explanation* (Oxford: Clarendon Press).
Millar, A. (2008). 'Perceptual-Recognitional Abilities and Perceptual Knowledge' in A. Haddock, and F. Macpherson (eds.) *Disjunctivism:Perception, Action, Knowledge* (Oxford: Oxford University Press), 330-47.
Millar, A. (2009). 'What Is It That Cognitive Abilities are Abilities to Do?' *Acta Analytica*, 24, 223-36.
Millar, A. (2010). 'Knowing From Being Told' in A. Haddock, A. Millar and D. Pritchard (eds.) *Social Epistemology* (Oxford: Oxford University Press), 175-93.
Pritchard, D., Millar, A., and Haddock, A. (2010) *The Nature and Value of Knowledge: Three Investigations* (Oxford: Oxford University Press).
Ryle, G. (1963) *The Concept of Mind* (Harmondsworth, Middlesex: Penguin Books).
Searle, J. (1969). *Speech Acts: An Essay in the Philosophy of Language* (Cambridge: Cambridge University Press).
Sellars, W. (1963). 'Some Reflections on Language Games', in Sellars's *Science, Perception and Reality* (London: Routledge and Kegan Paul), 321-58.

Alan Millar
Professor of Philosophy
Department of Philosophy
University of Stirling
Stirling
Scotland, United Kingdom

