

Throwing out the Tacit Rule Book: Learning and Practices

The definitive version of *Throwing out the Tacit Rule Book*, 2001, has been published in *The Practice Turn in Contemporary Theory*, edited by Ted Schatski, Karin Knorr-Cetina and Eike von Savigny, London: Routledge, 120-30, all rights reserved.

Stephen Turner
Department of Philosophy
University of South Florida

Abstract

The idea that the explanation of a social practice requires some sort of appeal to shared tacit rules is discussed in the light of recent connectionist accounts of the learning of grammar. The appeal of "tacit rule book" explanations in the case of linguistic regularities has rested on the fact that there seem to be no alternatives; similarly for accounts of social regularities. But in the case of grammar, it seems, it is possible to account for something like the learning of rules by reference to the external features of the information that is available. In the case of grammar, for example, it is the way information is structured, and its quantity, that enables individuals to generate consistent responses that resemble rule following. Social practices also can be thought of in terms of the way they are learned, and in terms of the information on the basis of which they are learned. This suggests some very elementary considerations about the way in which people learn what they need to learn in order to perform some activity, such as flirting or choosing clothes for the beach. What they need to learn is relative to individual purposes, which are diverse; what they experience, what information they get, depends in large part on these diverse purposes. Yet the activities are orderly. So it is implausible to conclude that the relevant tacit knowledge in these cases is knowledge of a common tacit rule book. No such book is required to explain the facts of order, and the way people learn suggests that they are not learning common rules.

"Practices" talk, I have argued elsewhere, gets into trouble over the notion of "sharing" (1994). The idea that there are "shared" practices requires some sort of notion of how they come to be shared, and this notion in turn dictates how practices can be conceived. If we decide that these difficulties are insurmountable, I argued, we can dispense with the notion of sharing altogether. Practices without sharing, to use a phrase favored in the nineteenth century, are habits--individual rather than shared. Habits are simply the part of the phenomenon described by the term "practices" that remains when the idea of people possessing the same shared thing is eliminated. "Habits," however, is a potentially misleading term, especially if "habit" is thought of as a generic alternative explanation rather than simply as the residue of the concept of practices once its objectionable elements have been eliminated. In what follows I will try to avoid this potential misunderstanding by restating my argument against the "social" conception

of practices in somewhat different terms, without appealing to “habit” as a concept, and by locating the argument in relation to recent work in cognitive science.

“Practices,” for the sake of the following, is defined as those non-linguistic conditions for an activity that are *learned*. By “a practice” I will mean an activity that requires its genuine participants to have learned something of this tacit sort in order to perform. What I intend to discuss are some general features of learning that constrain our conception of practices and therefore of a practice which depends on them. Ordinarily, the tacit stuff is not all there is to a practice. Most cases of a practice involve explicit communication or even explicit rules. Rules are not self-applying, so in the case where there are explicit rules, such as the law, the relevant practices are the practices that enable a person to follow the rules, for a lawyer or judge to interpret the law, for example. Sometimes there are no explicit rules, but there is explicit discussion. Painting a house, for example, can be done correctly or incorrectly, and there is a fairly elaborate vocabulary of evaluation and description of mistakes. Some kinds of “knowing how” that might be called “a practice” may have no such elaborate vocabulary of appraisal, and perhaps may have none at all. The practice of flirting, for example, before it was theorized about, presumably lacked such a vocabulary, and small children who flirt presumably have no vocabulary with which to discuss it-- but it nevertheless has to be learned. My concern throughout will be with the tacit parts of a practice.

The Linguistic Analogy

Language has always exercised a regulatory role in discussions of practice. Any account of practice that fails to account for language will be defective, because linguistic practices are part and parcel of many other practices and because linguistic practices are in principle not sufficiently different from other practices to regard them as likely to have a radically different character. The usual understanding of what is involved in the case of language is this: we communicate by virtue of sharing in the possession of this highly structured whole, a language, including the non-linguistic learned conditions for the use of the language, the practices. This notion can be put in a much more cautious way, as for example Davidson does when he speaks of “sharing a language, in whatever sense this is required for communication” (1987:166). The “required sense” of sharing may be minimal, and may not consist of shared tacit rules. In what follows, I propose to deal with the question of what the “required sense” is, and how it can be squared with a plausible account of learning.

Davidson’s remark is fairly conventional stuff in contemporary philosophy, but the argument that informs it is elusive. Is this a kind of unformulated transcendental argument, which amounts to the claim that the “sharing” of “language,” in some unspecified sense of these terms, is a condition of the possibility of “communication” in some unspecified sense of this term? Or is it a kind of inference to the best explanation in which there are no real alternatives-- an inference, so to speak, to the only explanation (which is perhaps not a bad definition of transcendental argument)? There are good reasons to be suspicious of arguments of this form. Yet this general picture, of some sort of shared (and presumably tacit) stuff at the basis of language, is highly appealing, and so is its extension to practices generally. The claim that there is some class of things that couldn’t happen, were it not for the existence of some sort of shared practices, is a commonplace, despite, and perhaps because of its vagueness.

Symbolic and Connectionist Models of Higher Cognitive Processes

As I have said, the Achilles heel of transcendental arguments is that the unique explanation to which the *explanans* point may not be the only explanation. In this case the argument establishes nothing. There is a close analogue to this kind of argument in cognitive science, and it has recently succumbed, at least in the view of many, to the demonstration that an alternative explanation suffices. The argument is this. People have the capacity to reason mathematically and speak grammatically. We can represent mathematical reasoning and the grammatical structure of a language explicitly, in terms of formal proofs and grammatical rules respectively. The fact that people can do in their head what can be done by formal proofs or in accordance with grammatical rules is a fact of the same kind as communication. It is the sort of fact that seems to require that the people who reason mathematically or speak grammatically possess capacities which pretty closely resemble, and operate like, those of formal proof. In short, people, in thinking mathematically or speaking grammatically, must be employing some sort of mental analogue to the rules of inference and axioms that go into mathematical proofs. The problem for the cognitive theorist is to model these capacities by identifying the tacit rules and axioms that are employed.

In cognitive science, this problem leads to a specific difficulty, the “central paradox of cognition,” stated by Smolensky, Legendre, and Miyata as follows:

Formal theories of logical reasoning, grammar, and other higher mental faculties compel us to think of the *mind* as a machine for rule based manipulation of structured arrays of *symbols*. What we know of the *brain* compels us to think of human information processing in terms of manipulation of a large set of *numbers*, the activity levels of interconnected neurons. Finally, the richness of human *behavior*, both in everyday environments and in the controlled environments of the psychological laboratory, seems to defy rule-based description, displaying strong sensitivity to subtle statistical factors in experience as well as to structural properties of information (1993: 382).

In this case there is an alternative explanation or approach, namely connectionism.

“Connectionism” refers to the claim that the appropriate model for the computation that occurs in the brain is not, as a once dominant viewpoint had it, the operation of logic machines that process symbols, but rather is the parallel distributed processing that is used on a variety of actual computer applications (such as flight simulators) and requires very substantial computing power. The “symbolic processing” model worked as follows: the mind acquires, either by genetic pre-programming or learning, rules for processing inputs, in a way that is familiar from ordinary computing, in which symbols come in well-defined forms and the computer program operates as if computational “rules” are “applied” to them mechanically to produce predictable outputs. Connectionist models work differently. The computer is given a learning algorithm, but no detailed “rules.” The computer is then “trained-up” by feeding it data and then giving feedback for “correct” answers. This is very much a Humean rather than a Kantian machine. Everything that is inside, except for the most basic capacity for forming “expectations” is a result of inputs, or experience. The inputs are not symbolic, but simply impulses originating from

various sensory sources, which are distributed through the brain in pathways made up of “connections” that are formed statistically, by the association of impulses of one kind with impulses of another kind. These are modeled mathematically as “weightings” of the impulses, which travel from “node” or pathway link to “node” and which modify the link by passing through it, just as a person walking in the forest makes a path, increasing the likelihood of future impulses of a similar kind being distributed in a similar way. The changes in the likelihoods are “learning.” These computer methods actually work: this is a model based on actual computer achievements, in which parallel distributed processing systems learn to do such things as detect cancers by being trained entirely empirically with inputs of images and feedback for correct predictions. No theory is needed, and no rules are identified or used in this method. The processes are statistical, and the capacities and outputs of the computer depend on what has been fed to it in the form of data and feedback.

The problem for modelers attempting to deal with human cognition is whether this approach is capable of accounting for higher mental processes. The general explanatory problem is the question of “how...competence that is highly systematic, coherent, compositional, and productive” can be achieved with the specific kinds of “finite and fixed resources” that connectionism employs (Smolensky, Legendre, and Miyata, 1993: 383). The highly influential paper by these authors from which these quotations are taken presents some technical results that bear on this problem. Indeed, in the opinion of most cognitive scientists and philosophical observers, these results represent a decisive turning point in the resolution of the issues. Briefly, what Smolensky, Legendre, and Miyata establish is that a connectionist account can be given of certain kinds of grammatical rules previously thought to be impossible to account for without reference to internalized formal rules. Their strategy is to show how “a fully distributed pattern of numerical activities” of a connectionist kind can be “the functional near-equivalent of a symbolic structure” (1993: 382). That is, they show how something like “rules” can be the product of “learning” through the simple mechanisms of spreading activation employed by connectionist accounts of the brain. The key idea in their analysis is that there is a kind of purposive process which occurs “when the ... activation spreading process satisfies certain mathematical properties” (1993: 383), a process they call maximizing Harmony. The strategy of the paper is to make rule-acquisition, by which they really mean the acquisition of functional equivalents to a rule, into a special case of connectionist learning generally. This raises an obvious possibility: that practices too may be special cases of this kind, or, alternatively, that practices may be better understood not as a special case of the same kind, but in light of the general properties of connectionist learning. In what follows, I will suggest that the latter is the most plausible conclusion.

I take it that there are two main implications of interest to the study of social practices of connectionism generally, one flowing from the other. The first is that because learning starts from a system state in which some learning has already occurred, we would expect that learning is always a product of a transition between a state A and state B such that the transition between another state in another machine and a state Bⁿ which is functionally equivalent to B will be a transition from a different starting point to a different end point. In short, the mechanisms in question, even if they are governed by the same basic simple mechanisms are individuated by the history of the mechanism. And, in general, the simpler the mechanism and the longer the chains of links between simple mechanisms the greater the diversity produced by differences in, so to

speak, the training history. Like paths from one point in space to another, the connections in a net that produce the “same” competency may be different in structure.

The implication of this that bears on the theory of social practices or the idea of shared practices is that two individuals with an ability to perform the general kind of task may go about it in ways that are quite different on the level of neuro-cognitive description. Put more simply, if we throw out the idea that there is a rule book that people tacitly master in order to, say, communicate we also throw out the idea that there is some single thing that people must all have in order to communicate. The approach taken by Smolensky, Legendre, and Miyata modifies this implication, for it suggests that something very much like “the same” rules of grammar may result from the fact that there is, in effect, a common end point to the process of mastering a grammar, namely maximal Harmony. But the approach also raises the question of when and where this notion of functional equivalence is applicable or relevant.

It suggests the following answer to the question: when information is plentiful and structured in such a way that “optimizing Harmony” or some other quasi-purposive system goal can lead to the same rule-like results. The general point supported by connectionism is the idea that the simpler the mechanisms that are the building blocks, the longer or more complex the total structure producing the result will be. Unlike rule books, these various complex individual mental structures are built up over time on the basis of different learning events. Ordinarily there will be a significant diversifying effect: the individual facts of the history of the acquisition of many cognitive skills will differ in such a way that the results differ. The question is where the result is rule-like and the rule-like structures are shared, and where it is either not rule-like or not shared. I note, incidentally, that Smolensky, Legendre, and Miyata say nothing about the question of whether more than one individually generated reduction of complexity can be functionally equivalent to a grammatical rule. However, there seems to be no reason that more than one result might be optimally Harmonious. Moreover, actual speakers do vary, so there is no reason to think that there is even one set of rule-equivalents that the process of harmonizing would necessarily tend toward in any given language.

We cannot, of course, answer questions about the existence of shared functional rule-equivalents directly. But some light can be shed on them by considering the ways in which “rules” are learned. Consider the child’s acquisition of the ability to perform simple arithmetical tasks. What is it to “be able” to add $2 + 2$? Is it merely to parrot the correct answer? Presumably not. Indeed, there may be no “criteria” in a Wittgensteinian sense for the possession of this competence. Obviously, a child does not master arithmetic immediately or all at once in the sense that the capacity is turned on like a switch. Other things must be mastered first, like counting, and these are often things that it is quite unproblematical to suggest are mastered in different ways. Some children may count on their fingers. Others may learn through singing the numbers, and others may master a great deal of material by rote without knitting it together, and only later make connections between the numbers of a mathematical kind. In short, students come to the learning of $2 \times 2 = 4$ from different starting points. They are then put through a series of experiences, and of course each student experience is slightly different and each classroom's experience is different. The differences, however, are not supposed to make a difference in the performance of the capacity. There are right answers, and the point of the various experiences with students having various prior experiences is that the experiences taken together transform the child cognitively in such a way that the child is able to perform the

cognitive task correctly. Almost everybody manages to do this. At the period prior to mastery, the cognitive architecture which supports the child's efforts will be, according to the picture I have given here, different. Different children will have different experiences on the way to mastery and the cognitive architecture will be a product of the path and the experiences along this path that the child takes from its starting point to the goal of mastery of the cognitive task. The purposes of children will vary as well. There may be a complex heterogeneity with respect to the goals. Some children may wish to avoid the embarrassment of being brought before the blackboard and humiliated for making mistakes. Other children may have a more positive experience of mastery and pride in achievement. These differences, like differences in the history of learning, do not have any effect on the competence itself.

The question is why? On the account of mastery that fits best with the tacit rule book model, the reason for this is essentially as follows: mastery is no more and no less than “getting” the basic rules of arithmetic, which are the same for everyone. The student tries this and that, gets corrected, gets told the answer is correct, and does all of this without understanding. But at some point something clicks, and the student “has” the rule. The history of acquisition is irrelevant because the important moment is the moment of clicking on to the rule, of getting the rule. This model of learning, what I will call the snap-on model, makes the history irrelevant. There is a radical difference of kind between the period before acquiring the rule and the period after, and the learning events of the first period have no effects in the second.

This is obviously an appealing story. It fits well with a certain view of Wittgenstein, and indeed may-- though I doubt it-- represent the most plausible explication of his notion of rule-following. I do not wish to take the issue up here directly, but I will note that in the history of the reception of the *Philosophical Investigations* there was a period in which something like the account I have given here was purveyed by Wittgenstein's students and interpreters.¹ Nevertheless I think it misleads us about practices generally, misleads us into looking for “criteria” or “agreements” where there is nothing of the sort to be found. The main reason for this is that mastery-- however one wishes to think of it-- is in most cases *not* the same thing for different people under different circumstances. It is purpose relative, and the purposes of individuals involved in the activity vary. It is also situation or experience relative, in the sense that it depend on the materials to which the rules are applied.

Differences in purposes lead to differences in experience, and this means differences in the information that is fed into the system. Diversity is the normal result, but diversity is nevertheless consistent with a great many kinds of cooperation, and indeed, I think, with communication. In what follows I will consider some examples, and suggest that most of those we call practices are more plausibly thought of as the common activities of people with diverse learnings than as activities made possible by the sharing of the same rule-like structures. Obviously there is no room here for knock-down arguments. Indeed, the complexity of the processes involved assure, I think, that they will always be opaque to analysis. But something may usefully be suggested about the probable effects of differences in purposes on the cognitive side of practices.

Notice a few unusual features of instruction in arithmetic. Children are tested on their mastery of multiplication tables and there is clearly a right and a wrong answer to such questions. Children are then disciplined or drilled in these right answers, and indeed in days gone by, the times tables were simply mastered by rote and no attempt was made at giving the

child some sort of conceptual understanding of multiplication. This kind of training is anomalous if we consider the universe of “practices.” Whatever this universe might be taken to consist of, presumably it includes such “practices” as the standards of etiquette that Norbert Elias describes, the habits of moderation and compromise that are the means of assuring the fruitfulness of parliamentary discussion that thinkers from John Austin to Michael Polanyi and Michael Oakeshott have supposed to operate at the heart of British parliamentary politics, and perhaps many other things as well, such as the means by which laboratory scientists identify objects, as well as the examples I have given here of flirting and house-painting.

Does the snap-on account of learning fit these cases as well? One difference is this: in these more complex cases, perhaps with the exception of flirting, there is a large spoken or explicit element, a vocabulary of appraisal or even theory about the activity, in addition to its tacit base. In the case of elementary arithmetic there are theories, but they are known to mathematicians alone and are not part of the activities of making change and the like in which elementary arithmetic is used. Another difference with the case of arithmetic is that in these “complex” cases the means by which information about right and wrong are conveyed is different, and the purposes that the parties to the practice have are diverse and lack a common core, like making correct change. So it would be a bit surprising if a model in which these differences had no place made much sense of the learning of the practice. If we see the teaching of arithmetic as the employment of various behavioral technologies that are designed to produce absolute consistency of response, it is no surprise that such things as the differences in purpose between learners have no effect on what is learned-- that is one of the incidental consequences of a behavioral technology that is designed to eliminate differences. But the behavioral technologies that serve to convey information about right and wrong with respect to these other bodies of practice do not work in this way.

Grammar is learned without behavioral technologies of the same explicit kind. But like arithmetic, the effects of diversity of purposes are overwhelmed, by the quantity of redundant information and, perhaps, by the structure of the information on which learning operates. Flirting, in contrast, is information-poor, and the starting points of individuals vary enormously. So it would be odd to find functionally equivalent rules there. Nevertheless, flirting, or at least a personal way of flirting and responding to flirting, is learned, and thus fits the model of a practice with which I began. Similarly for politics. One simply doesn't have the vast amount of experience necessary to overwhelm the diversifying effects of differences in experiences, difference in starting points and differences in purposes. The explicit rather than the tacit parts of politics, the vocabulary of appraisal, the body of political and historical discussion, and explicitly formulated beliefs of various kinds, do the work of making the practice hang together. A practice such as scientific discovery, built around training that is oriented to enabling a person to participate in discussions involving highly specialized terms and employing common apparatus, may in some respects be more like arithmetic, with its explicit behavioral technology of tests, at least with respect to the mastery of techniques. But scientific discussion itself often skates on the border of mutual intelligibility, and not infrequently goes beyond it. And explicit discussion, not the training base, pulls the practice in new directions and toward new goals and experiences.

Practices in the Sociological Sense: Social Order

The characteristic way in which sociologists have analyzed social practices in the past has been in terms of patterns which are observed, in which the analyst can say that people behave as if they are following a particular rule, and in which the analyst can point to some sort of sanction or response to violations of the “as if” rule that indicates that some behavior is deviant. Expressing disapproval is a much different means of conveying information than correcting arithmetic tests. Take the practice of dressing for the beach. Beach-dressing practices differ from country to country and place to place. They differ more or less systematically. What is appropriate in one place, or for one sort of person, is not appropriate for another. There is no place in which one can look up these “rules.” One may be entitled, from this, to conclude that there is among beach-goers in particular places some sort of tacit “code” which forbids certain kinds of attire or defines appropriate attire. But this is very peculiar sort of conclusion. It seems to be little more than a short hand for saying that if one does various things, some people will express disapproval. One can get quite an elaborate account of the whole business of approval and disapproval, of the distinctions that are implicit in the pattern of disapproval and approval, and so forth.

But it is far less clear that we are all attempting to master the same tacit rule book, or indeed that there is some sort of common thing that is being mastered here. Are we simply trying to dress for the beach in accordance with various purposes, one of which may be-- but may not be-- to avoid disapproval? Could it be that the apparent structure of the activity is simply the results of different people with different purposes and a variety of attitudes acting and expressing approval and disapproval? The tacit rule book model requires us to think that, as with arithmetic exams, we face social life as an exam which we respond to in terms of approval and disapproval. If we meet with approval, we may take it as a sign that we have not violated the tacit rule book. Disapproval represents failure in mastering the tacit rule book. We learn from the experience what is in the tacit rule book or how to apply the tacit rules. If we retain the image that what we have here is a vast tacit rule book of great complexity then we can think of individuals as having, though not very satisfactorily, mastered elements of or approximations of this tacit rule book, that is to say as having incomplete or unsatisfactory mastery.

But this seems rather strange, for a variety of reasons. Go back to the child's performance of addition in a schoolroom. The child's goal may be for the most part to avoid embarrassment. Getting the right answer is simply a means to this and doubtless other ends. Satisfactory performance from the child's point of view is one that manages to achieve enough of the things the child wants to achieve or to avoid the things the child wants to avoid. Of course, children manage to satisfy in the classroom in lots of ways, such as by deflecting embarrassment by becoming the class clown. In short, because goals are heterogenous, "mastering" is also a heterogenous notion.

With the usual sorts of things we think of as social practices, such as dressing to go to the beach, our goals are heterogenous. What counts as satisfactory will be different for different people. And this has implications for the experiences the person has, and consequently the information that they receive as a result of these experiences. The beach-goer who avoids arrest who is satisfied with the knowledge that enables him or her to do this is different from the beach-goer who is satisfied with admiring glances or with being ignored and thus avoiding social disapprobation.

Why should this make any difference to the question of whether there is a tacit rule book

of beach attire, or flirting? Arithmetic problems are odd in that optimal mastery is necessary with respect to the large number of usual goals, such as avoiding humiliation, that children bring to the situation of learning elementary arithmetic. Ordinarily matters are different: different goals dictate different results, because the pursuit of the goals leads to experiences that are different in kind. Sometimes, however, there are intermediate goals that are the same, or demand functionally equivalent conduct or understandings. Language seems to be a case like this. Intelligibility is an intermediate goal that is shared by a large number of people. The mastery of a language is necessary for the achievement of this goal. People have very diverse linguistic experiences. But the quantity of linguistic data with which they operate is so immense that it obliterates many of the differences. In the case of beach attire, matters are quite different. It is simply an overstatement of the case to say that people have to internalize a norm of dress in order to participate in the activity. There is no single “norm” that corresponds to the various “mastering” strategies that people exhibit in the course of responding to the problem of appropriate beach attire. Not only do people have different explicit ideas about what is appropriate and what is not, they *respond* differently. And their different experiences lead them to have quite different kinds of information on which to operate, information that is different in kind, and not redundant and overwhelming in such a way as to produce the same results. Yet there will be something recognizable to us, as analysts, as a social practice.

If this is a possible explanation, it is a better explanation than the idea of a tacit rule book. The experience of social life, in the absence of a massive amount of information or of highly structured experiences, such as the learning of arithmetic, is simply too diverse, and too thin, for the individual to derive from it anything so determinate as a set of rules that is the same for everyone. Even if there were such rules, and even if individuals optimized in acquiring functional equivalent rule-like responses to the material presented to them, it is reasonable to suppose that the material on which each individual works is too limited-- in all but the exceptional cases-- for them to acquire the same rule-like response as the next person does.

The Domain of Practices

If people are “master” learners who start from different points and acquire what they learn through different sets of experiences and who satisfice according to different goals which may change over time and thus direct the path of experiences and learnings in different ways, it may appear that the real mystery here is how there could be any such things as social practices or “social order” at all unless there is massively redundant information structured in the appropriate way to produce functional equivalents of the same rule in all parties to the practice. I think that this is indeed the right question and the answer to this question needs to be not that there is a tacit rule book and that the problem is to figure out how people acquire it but rather that the kinds of patterns and regularities we regard as social practices are nothing more than that which people learn, in a rather heterogeneous way, are the best ways or the satisfactory ways to negotiate the paths toward to the fulfillment of whatever purposes they might have. What this suggests is that what people acquire that sociologists call practices are lessons that enable them to do particular things, such as go to the beach and be comfortable with the responses of other people or at least get from them responses that they are satisfied with.

This gives us the basis for a crude taxonomy of practices. The learning of some practices

is indifferent to the purposes for which they are learned, or, rather, “optimizing Harmony” makes them the same for everyone. For other practices, in contrast, mastery is purpose-relative. Neither case, I have suggested here, *requires* the model of a common tacit rule book. It is mistaken in the cases in which mastery is purpose-relative because it is simply the imposition of the sociologist-- no one learns it, and there is no plausible way for a complex scheme of tacit rules of the hypothesized kind to be learned. It is mistaken in the case in which optimizing Harmony results in the sharing of rules that are functionally the same for everybody because it involves a mistaken inference from the explicit form of an activity to a supposed cognitive basis with a similar form.

References

- Davidson, D., “The Method of Truth in Metaphysics,” *Midwest Studies in Philosophy* II, 1977, pp.244-254.
- Smolensky, P., Legendre, G. and Miyata, Y. (1993) “Integrating Connectionist and Symbolic Computation for the Theory of Language,” *Current Science*, 1993, vol. 64, pp.381-391.
- Turner, S., *The Social Theory of Practices*, 1994, Oxford, Polity.
- Winch, P., *The Idea of a Social Science*, 1958, London, Routledge.

Notes

1. "We see that we understand one another, without noticing whether our reactions tally or not. *Because* we agree in our reactions, it is possible for me to tell you something, and it is possible for you to teach me something," as Rhees puts it (quoted in Winch, 1958: 85). This is impeccable Wittgenstein (cf. 1979: 89), up to a point-- the point at which agreeing in our reactions is made into something like a criteria of the existence of a rule.