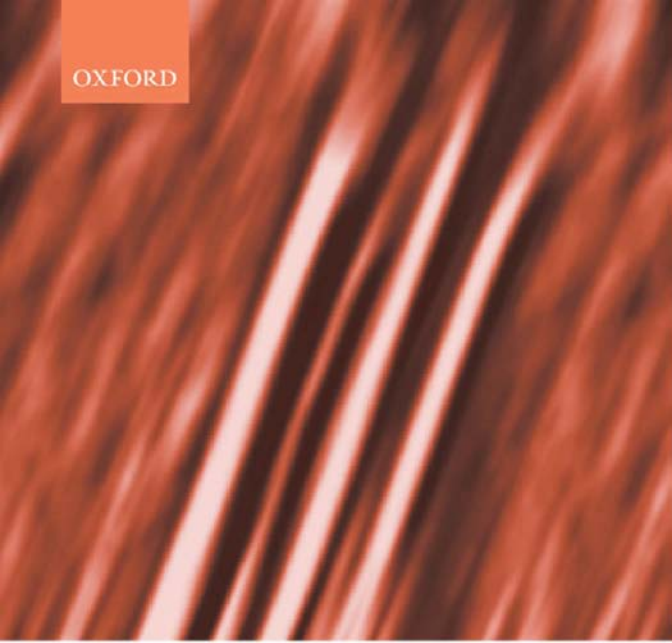

Knowledge by Agreement

MARTIN KUSCH

Contents

Chapter 14 NORMATIVITY AND COMMUNITY

NORMATIVITY, THE PRIVATE LANGUAGE ARGUMENT, AND RULE-FOLLOWING

One of the most central premisses of communitarian epistemology is the idea that normative phenomena—rules, norms, conventions, prescriptions, and standards of correctness—can exist only within communities. Put differently, socially isolated individuals are unable to generate normative phenomena. Let me call this idea the 'community thesis'. In what follows, I shall try to defend it.

In contemporary philosophy the community thesis is being discussed primarily in the philosophy of language and mind. More precisely, the community thesis is crucial in arguments against so-called 'private languages'—languages that are used only by a single, socially isolated, individual. Philosophers seeking to demonstrate the impossibility of private languages rely on the link between normativity and community as a premiss in their overall argument. Put in a nutshell, this argument goes as follows:

For a system of signs to count as a language, actions of sign-use must be classifiable along two related dimensions: the dimension of 'correct versus incorrect', and the dimension of 'seems right versus is right'.

A phenomenon is normative if it essentially involves the distinctions 'correct versus incorrect' and 'seems right versus is right'.

Normativity presupposes community.

Ergo: Private languages are impossible.

Often the argument is conducted on a slightly more general level: the question then becomes whether a socially isolated individual would be able to follow rules (regardless of whether or not these rules are rules of a language). In this more general form the argument runs thus:

One cannot speak of an actor as a rule-follower unless the distinctions 'correct versus incorrect' and 'seems right versus is right' are applicable to the actor.

A phenomenon is normative if it essentially involves the distinctions 'correct versus incorrect' and 'seems right versus is right'.

Normativity presupposes community.

Ergo: Private rule-following is impossible.

Much of the dispute over the possibility of private rule-following is focused on the community thesis. 'Individualists' defend the possibility of private rule-following by denying the thesis; 'communitarians' oppose the possibility of private rule-following by insisting that the thesis is correct. I believe that the communitarians have won the argument, but in order to make this assessment plausible, I shall have to present some key junctures of the dispute.[78]

The debate over the community thesis has two interrelated levels. On one level, philosophers argue for and against the community thesis itself. On another level, and often in the very same paper, philosophers seek to show that Ludwig Wittgenstein is best interpreted as sharing their respective view of the community thesis. The two levels are almost always related to one another not only because the whole issue of private languages was first made salient in Wittgenstein's *Philosophical Investigations* but also because Wittgenstein is—rightly or wrongly—thought of as the ultimate authority in this matter. It seems unimaginable to many philosophers that Wittgenstein could possibly be on the wrong side of this fence. Here I shall nevertheless ignore the question where Wittgenstein himself stood on the issue. Although I happen to believe that he favoured the communitarian side, I shall not attempt to defend this claim in this book.[79]

[78] Needless to say, I am not striving for completeness here. Important anthologies on the philosophy of rule-following are Pitcher (1964) and Canfield (1986).

[79] As far as the interpretation of Wittgenstein is concerned, I have learned most from Bloor (1997), Canfield (1996), and Malcolm (1995).

DISTINCTIONS

In order to make the dispute over the community thesis tractable and decidable, it is important to give it a precise content. This is best done by making a number of distinctions.

1. We need to differentiate between two kinds of private language: those that are *essentially* private and those that are merely *accidentally* private. A language is essentially private if the meanings of its expressions are, or depend on, 'epistemically private items' (Craig 1982, 1997). The clearest examples of such items are sensations. You cannot have my sensations, and I cannot have yours. Hence, if our respective sensations are part of the meanings of our respective words, then our respective meanings must differ. Add to this the idea that one can never know someone else's sensations, and we arrive at the claim that all language is essentially private. A language that is merely *accidentally* private does not involve epistemically private items. Someone's language is accidentally private if de facto it is not shared with others—although it could in principle be shared. A secret language that I develop just for my own use would be accidentally private. Sometimes the term 'private language' is reserved for essentially private languages; accidentally private language is often called 'solitary language'. I shall follow this usage here. For my defence of the community thesis I shall only be concerned with solitary languages and solitary rule-following.

2. It is important to distinguish between *physical* and *social isolation* (Bloor 1997: 92; cf. Kripke 1982: 110). Social life is episodic; between our face-to-face encounters we are physically isolated from others. But even when we are physically isolated, we remain socially 'embedded': we are related to others in our thinking, feeling, and acting. We anticipate blame and reward, we make plans that involve others, and we harbour various feelings for them. To be physically but not socially isolated is a frequent experience for most of us. But we are never fully socially isolated. A wolf child (before capture) would be a case of someone who was both physically and socially isolated from other humans. And maybe a severely autistic child can be thought of as being socially but not physically isolated. The case of Defoe's Robinson Crusoe is not so clear. Crusoe grows up in human communities and washes up on the shores of his island only by the time he is already an adult. Subsequently he spends twenty-six years on his own. Was Crusoe socially or

just physically isolated? Intuitions diverge at this point. If we judge him to have been merely physically isolated, we are thinking of his twenty-six lonely years as a (very long) interval between face-to-face encounters. If we opt for social isolation, we are guided by the idea that twenty-six years is too long a period to count as an interval between social interactions.

3. We need to make a number of distinctions regarding rules. One important distinction is that between *rules* and *regularities* (Robinson 1992). Some regularities of human behaviour are 'natural' and innate. Instinctive behaviour fits this bill. Other regularities of behaviour are 'acquired'. Among the latter it is useful to distinguish between those that are acquired voluntarily and those that are acquired involuntarily. Conditioning is a case of the latter, learning a skill a case of the former. Not every regularity of human behaviour is due to a rule, even though the practice of every rule-follower will display some regularity. The natural regularities in an agent's behaviour obviously are never caused by the agent's following of rules. But neither are involuntarily acquired regularities. Rules can be deliberately violated; this precludes the possibility of conditioned behaviour being an instance of rule-following.

4.*Following a rule* differs from merely *acting in accordance with a rule*. Londoners using the Underground know the following rule-governing behaviour on the long escalators: 'stand right, walk left'. They learn the rule by reading the signs. Many tourists in London do not understand English, and hence do not know of this rule; and yet most of them stand right anyway. They act in accordance with the rule but they do not follow it. A rule-follower must, on some level, know both what she *is* doing and what she *is supposed to* do.

5. It is customary to distinguish between the *introduction*—or *stipulation*—and the subsequent *application* of a rule. The introduction stipulates the content of the rule, what is to be done under what circumstances. When we apply rules, we judge whether the given circumstances are those specified by the rule. This distinction feels intuitive and simple, but it is important to recognize how misleading it can be. We cannot draw a sharp dividing line between application and stipulation of the content of a rule, since it is only in and through a series of applications that the rule acquires a clear content. Assume my wife and I introduce the rule that we go running in the morning. Surely, when first introducing the rule, we have not yet foreseen all the varied circumstances that might lead us to modify the rule: visits of relatives, illnesses, sports broadcasts from the Olympics in Australia, lack of time due

to early-morning appointments, hangovers, etc. And yet it seems natural to say that it is only in the process of our facing these varying circumstances of application that the rule itself acquires more and more content. Paradoxical as it might sound, it is only by applying the rule that we stipulate its content.

6. A further important distinction is that between *rule-formulation* and *rule-practice*. Rules can often be formulated in language: 'We should go running every morning', 'Most fish dishes call for white wine', 'Essays have to be handed in at the beginning of January'. But an agent can be said to follow a rule even if he is not thinking of the rule at all. On many mornings my wife and I will go running without even considering our rule. Or think of a baseball player (Searle 1995: 146). He would be greatly impeded in his performance if he had to think of the rules while chasing and throwing the ball. These scenarios are cases where a rule-practice has evolved. A rule-practice is a pattern of actions that the actors regard as being correctly captured by the rule-formulation. The practice persists and develops on the basis of precedent rather than rule-formulation. The actors do 'what we have always done' rather than 'what could be described as falling under the rule-formulation'. This is not to say that the rule-formulation is irrelevant. The possibility of discourse about the practice is what makes it a *rule*-practice rather than an *instinctual* practice. And citing the rule-formulation is an important part of that discourse.

7. The word 'rule' is 'an especially messy cluster concept' (Lewis 1969: 105). We speak of rules in many different contexts. Sometimes 'rule' is used to refer to generalizations or laws of nature, sometimes for maxims or hypothetical imperatives, sometimes for norms and conventions. A detailed analysis of the various uses would be important and helpful, but it has not to date been attempted. Nevertheless, some uses of 'rule' have received closer attention. Interestingly enough, these all tend to be uses that involve communities rather than individuals. In other words, analysts have focused on social rules. The perhaps best-known account was suggested in the early 1960s by the legal theorist H. L. A. Hart; it has recently been amended and improved by the social philosopher Margaret Gilbert. Gilbert offers the following definition of a social rule:

There is a *social rule in a group G* that action A is to be done in circumstances C, if every member of G:

[*a*] regularly does A in C…
[*b*] has a 'critical reflective attitude' to the pattern of behavior: doing A in C…
[*c*] regards doing A in C as a 'standard of criticism' for the behavior of members of G…
[*d*] regards non-performance of A in C by a member of G as a fault open to criticism…
[*e*] criticizes any member of G who does not do A in C and puts pressure to conform on members of G who threaten not to do A in C…
[*f*] believes that every group member *has a claim against* every other group member for the performance of A in C and a consequent *title to exert punitive pressure* on any other group member in favor of doing A in C…
[*g*] expresses their criticisms and demand using normative language such as 'You *ought* to do A now!' or 'That's wrong'…
[*h*] feels that members of G are in some sense 'bound' to conform to the pattern: doing A in circumstances C…(Gilbert 2000: 73, 76)

I find this analysis adequate as an account of social rules. The question we will have to ask ourselves below is how much we can give up of this deeply social picture, while retaining the notion of a rule.

8. As we shall see below, all participants in the debate over the community thesis agree that some level of 'multi-subjectivity' is needed for there to be rule-following. In other words, all sides agree on this point: an agent who exists only for one single moment in time could not be said to follow a rule. The disagreement begins after this point and concerns the question whether rule-following presupposes *intersubjectivity* or mere *intrasubjectivity*. Does rule-following presuppose a community of different individuals or will a community of the time-slices (or modules) of one and the same individual do? Needless to say, these two positions will only mark a clear distinction if we ignore multiple personalities, time travel, and people with memory disorders.

9. It is useful to reflect on the kind, or strength, of argument one expects to emerge from the debate over rule-following. Many arguments in the debate are appeals to semantic intuitions concerning the word 'rule'. Communitarians insist that 'rule' really means social rule; individualists are adamant that the category 'rule' is not exhausted by social rules. The difference comes out clearly in disagreement over whether a congenitally isolated individual could invent rules for himself. Communitarians feel that speaking of the social isolate as following rules is stretching the concept of

rule too far. Individualists propose to see the social isolate as analogous to the social group. Communitarians reply that analogy is not good enough. Individualists think this reply begs the question. And so on. Unsurprisingly, not everyone is happy with the expectation that proper concentration on what we *really* mean by 'rule' will decide the issue. More than one philosopher has voiced doubts about the project of a conceptual analysis that involves checking our semantic intuitions in highly unusual circumstances. Leslie Stevenson puts the point as follows:

> Why should we expend so much mental effort pursuing the question of the mere logical possibility of these shadowy creatures of myth and philosophical fantasy—Adam before the creation of Eve, a Crusoe isolated from birth, the articulate wolf-child, the solipsist with no concept of other minds? All the evidence goes to show…their psychological impossibility…Since actual cases of articulate Crusoes have not come up, we have not had to decide what to say about them. And, on reflection, we can see that there is no reason why our concepts of meaning, knowledge, rule-following, etc., should somehow have built-in guidance for how to apply them to cases never before encountered. (1993: 447)

I sympathize with the general sentiment but I am somewhat less pessimistic. The communitarian viewpoint is strengthened if it can be shown that the individualist's case rests on highly implausible interpretations of artificial scenarios. We should not expect it to be an analytic truth that rule-following is social. But we might come to see that the social nature of rule-following is central to how we think about language, normativity, and ourselves.

10. Finally, the community-dependence of rule-following is not a simple 'yes or no' issue. Community dependence is a matter of degree. A weak degree is acceptable even to the individualist; and communitarianism can be more or less strong. We need then a taxonomy of degrees of community-dependence. Such taxonomy was first suggested by Colin McGinn (1984: 194–9); mine is a further development of his.

(A) *Strongest Present-Tense Community Thesis.* An individual is able to follow a rule only if the individual is currently a participating member of a group in which the very same rule is followed by other members.

(B) *Strong Present-Tense Community Thesis.* An individual is able to follow a rule only if the individual is currently a participating member of a

group in which some rules or other—but not necessarily the same rule—is followed by other members.

(C) *Past-Tense Community Thesis.* An individual is able to follow a rule only if the individual *has been, at some point in the past,* a participating member of a group in which the very same rule, or some other rule, was followed by other members.

(D) *Moderate Community Thesis.* An individual follows a rule if and only if the participating members of an existing group judge parts of the individual's behaviour to be *similar to* their behaviour in following the very same rule. The individual in question need not be a member of the group in question.

(E) *Weak Community Thesis.* An individual follows a rule if and only if it is possible to imagine a community in which that rule is followed by participating members.

(F) *Weakest Community Thesis.* An individual follows a rule if and only if the participating members of an existing group judge aspects of the individual's behaviour to be *similar to* aspects of their behaviour in following some rule or other. The individual in question need not be a member of the group in question.

I shall refer to A to E as communitarian views, and to F as the individualist position. Eventually, I shall defend A, a position also advocated in Bloor (1997). C is the most commonly defended communitarian position (e.g. Malcolm 1995; Meredith Williams 1991). D has been put forward by Kripke (1982: 110). E is favoured by Peacocke (1981: 93–4). And F is put forward by, among many others, Baker and Hacker (1984). Most individualist attacks on the community view have been directed (jointly) at A, B, and C. D and E have been scrutinized less often.

AGAINST WEAK AND MODERATE COMMUNITY THESIS(D AND E)

Peacocke states the *Weak Community Thesis* E in the following passage:

> Nothing in the community view…excludes the possibility of a permanently isolated desert-islander rule-follower. The community view can count such a person

as a genuine rule-follower if he reacts to new examples in the same way as would members of our own community, or of some other conceivable community. The community need not be that (if any) of the rule-follower himself. (1981: 93–4)

The problem with this weak version of the community view is that it places too few constraints on rule-following: 'the subjunctive, together with the requirement of mere *conceivability* of a community, robs the qualification of any restrictive content whatever' (Baker and Hacker 1984: 449; cf. Budd 1984: 319). Moreover, one might ask whether Peacocke's position in this passage actually fits with the view expressed in the rest of his paper. A few pages before the quoted passage Peacocke summarizes Wittgenstein's—and his own—view as follows:

According to Wittgenstein, it is, roughly speaking, the way that others go on in new cases in applying a word, and my conformity with their practice, which determines whether I am following a rule, whether I mean something by an expression or not; and if this is the account; then it is not clear how someone could still mean the same as the other members of the community by some expression and yet try to revise the judgments which, according to Wittgenstein, help to determine its meaning. (1981: 89)

Revising community judgements, conforming to the practice of others—these seem activities that require an actual rather than a merely possible community.

The *Moderate Community Thesis* D is part of the most controversial interpretation of Wittgenstein's later philosophy, Saul Kripke's *Wittgenstein on Rules and Private Language* (1982). In a famous passage of this book Kripke denies that his communitarian interpretation of Wittgenstein makes the latter an opponent of all forms of private rule-following:

if we think of Crusoe as following rules, we are taking him into our community and applying our criteria for rule following to him. The falsity of the private model need not mean that a *physically isolated* individual cannot be said to follow rules; rather that an individual, *considered in isolation* (whether or not he is physically isolated), cannot be said to do so. …Our community can assert of any individual that he follows a rule if he passes the tests for rule following applied to any member of the community. (1982: 110)

Kripke's view seems to be that even extreme physical isolation (perhaps even physical isolation from birth) need not give rise to social isolation and thus

to the inability to follow rules. Crusoe's social isolation is avoided only if we find ourselves able to attribute to Crusoe the intention to follow one of our rules. For this Crusoe needs to fulfil the criteria for rule-following that we routinely apply to each other in our community.

This strikes me as an incoherent view. It clearly has the consequence that someone might be called a rule-follower even though he never interacts socially with any of us. And that seems to fly in the face of our very criteria for picking out rule-followers among us. When we try to decide whether one of us is following a specific rule, we inevitably engage in social (discursive) interaction. Indeed, the ability to explain oneself is, for us, the central criterion for rule-following. If we stick to our criteria, then Crusoe cannot be regarded as a rule-follower.

A different way of making the same point is to distinguish between Crusoe's and our own categories. For us to be convinced that Crusoe is a rule-follower, *he* must be able to draw a meaningful distinction—'meaningful' by our lights—between 'seems right' and 'is right'. And given our standards for attributing rule-following to an actor, we could not be so convinced unless we were to interact with him discursively. Needless to say, we could of course describe Crusoe's actions in light of *our* distinction between 'seems right' and 'is right'. But the possibility of this description is not enough for Crusoe to pass as a rule-follower; this description merely makes him act in accordance with a rule. (Note that I am not (yet) insisting that Crusoe could not follow rules. I am merely arguing that he could not be said to follow rules if such attribution was dependent on him fulfilling the very same criteria that we in society routinely apply to one another.)

Individualists' criticism of Kripke's moderate community view obviously proceeds differently. Individualists insist—in accordance with F—that we can attribute rule-following to Crusoe even if he does not abide by our rules:

> Must Crusoe's rules be the same as ours?…Could he not invent new rules, play new games? To be sure, in order to *grasp* them, we must understand what counts, in Crusoe's *practice,* as following the rules. And that must be evident in Crusoe's *activities.* But that is not the same as checking to see whether his responses agree with ours, let alone a matter of 'taking him into our community'. And our judgment that he is following his rules is quite independent of any judgment about how most members of the English Speaking People would react. (Baker and Hacker 1984: 434)

I find view F unconvincing. But in order to make my case as best I can, I had better first introduce the individualists' arsenal of arguments against stronger community theses A to C.

THE INDIVIDUALISTS' CRITIQUE OF STRONG COMMUNITY THESES(A TO C)

In order to be comprehensive and fair in my treatment of the individualists, I shall introduce all six of their major arguments against A to C. Following this, I shall rebut each of these arguments, often relying on help from other communitarians.

I propose calling the first argument the *Metaphysical Argument*. It appeals to our intuitions of what is conceivable, and hence metaphysically possible. The thought is this. There is nothing inconceivable about a scenario in which God creates a single rule-following individual. And since this scenario is conceivable, it is metaphysically possible. The inventor and champion of this line of thought is Colin McGinn. Here is his formulation:

> A vivid way to appreciate the force of what I am calling the natural view is to imagine what it would take to create (in a God-like way) rule-followers or concept-possessors. According to the community view, nothing you can do in the constitution of an individual can justify the assertion that he has concepts or means anything *until* you create other individuals…But how could the creation of other individuals wreak such a momentous change in the first created individual? (1984: 191)

The last sentence of the quotation foreshadows the second argument, the *Internalist Argument*. It insists that whether or not an individual is a rule-follower depends solely on facts about that individual's mental states. The point can be made forcefully by imagining catastrophes that have but a single survivor: 'could not the rest of the human race be wiped out while you sleep and yet the next day you awake with your rule-following capacities intact?' (McGinn 1984: 198).

McGinn is also the inventor of the third anti-communitarian line of thought, the *Invention Argument*. He uses it specifically against A but it could also be used against versions of B and C. The idea is simple. If all rule-following had to be communal, then rules could never have a first follower.

But rules must have first followers, or else rules could never be invented. The inventor of a rule is its first follower (McGinn 1984: 195).

A fourth line of reasoning is a parity argument. Since two different parity arguments have been used by individualists, I call the fourth line *Parity Argument 1*. It is an attack on one of the key ideas behind the communitarian position. Communitarians often reason as follows. One can only properly be said to follow a rule if the distinction between 'thinking one is following the rule (correctly)' and 'following the rule correctly' is somehow applicable to the rule-follower. More precisely, for a speaker to use a sign correctly, three conditions have to be met:

(i) There must be something a speaker, *a,* can do that counts as an instance of 'thinking that he is using a sign correctly';
(ii) There must be something a speaker, *a,* can do that counts as an instance of 'using a sign correctly';
(iii) It must be possible for *a* to do what counts as an instance of 'thinking he is using a sign correctly' without at the same time doing what counts as an instance of 'using a sign correctly'. (Armstrong 1984: 51)

Communitarians insist that (i) to (iii) cannot be met as long as we consider a single speaker (or rule-follower more generally). Only the views of others—views as to how the rule needs to be followed in a given case—can provide a standard by which the individual's aiming for correctness can succeed or fail.

Parity Argument 1 seeks to unsettle this line of thought by insisting on a parity between individual and community. Individual and community are on par: if individuals cannot draw the distinction between 'seems right' and 'is right', then neither can communities as a whole. The individual cannot get access to an 'is right' by referring to his community, for his community can at best offer him a communal 'seems right'. Put differently, Parity Argument 1 poses a challenge to the communitarian: show me how it is possible for a whole community to do what counts as an instance of 'thinking it is using a sign correctly' without at the same time doing what counts as an instance of 'using a sign correctly'. As John McDowell puts the point:

The trouble is that there is a precise parallel between the community's supposed grasp of the patterns that it has communally committed itself to and the individual's supposed grasp of his idiolectic commitments. Whatever applications of an expression secure communal approval, just those applications will seem to the

community to conform with its understanding of the expression.…One would like to say: whatever is going to seem right to *us* is right. And that only means that here we can't talk about 'right'. (McDowell 1993: 261)

Parity Argument 2 insists on the parity of individual and group in a different way. The thought is that whatever social *inter*subjective structure the communitarian insists on, the individualist can always point to a similar *intra*subjective structure that can do the same work. For instance, if rule-following presupposes more than a momentary isolated individual, why go all the way to full intersubjectivity? Why not try to capture the necessary 'multi-subjective' check by means of time-slices of the same individual? Parity Argument 2 has been pressed most convincingly by Simon Blackburn:

The members of a community stand to each other as the momentary time-slices of an individual do.…And when the community says 'well, we just see ourselves as agreeing (dignify, compliment ourselves as comprehending the same rule)' the individual just borrows the trick, and compliments himself on his rapport with his previous times. (Blackburn 1984*b*: 294–5)

Parity Argument 2 has been developed further by Paul Gerrans (1998), who attempts to give a more detailed psychological account on how 'intersubjectivity within a single individual' is possible. As Gerrans sees it, we as individuals can form a community within ourselves by running simulations on our own actions. I am to imagine what I would do under given circumstances, and then count each such possible me as a separate individual.

Perhaps the favourite weapon in the individualists' arsenal is the *Similarity Argument*. It is meant to defend the idea that we can easily imagine (congenital) social isolates to behave in ways that leave us with no alternative but to attribute rule-following to them. More concretely, the behaviour of social isolates might lead us to ascribe to them actions of rule-introduction, self-correction in the light of a rule, and awareness of the 'seems right/is right' distinction. The Similarity Argument supports this idea with various thought-experiments. I shall mention the three most prominent here.

The first example was introduced by Baker and Hacker (1984: 433):

there is no reason why Crusoe should not follow a pattern or paradigm, making occasional mistakes perhaps, and occasionally (but maybe not always) noticing and

correcting his mistakes. That he is following a rule will show itself in the manner in which he uses the formulation of the rule as a canon or norm of correctness. Hence, to take a simple example, he might use the pattern—…—…as a rule or pattern to follow in decorating the walls of his house; when he notices four dots in a sequence he manifests annoyance with himself. He carefully goes back and rubs one out, and perhaps checks carefully adjacent marks, comparing them with his 'master-pattern'. And so on. Of course, he is *not* merely following his 'inclinations', but rather following the rule. And it is his behaviour, including his corrective behaviour, which shows both that he is following the rule, and *what he counts as following the rule*.

Several authors attribute the second example to Michael Dummett without, however, giving a page reference (Blackburn 1984*b*: 297; Gillett 1995: 194; Wright 1986: 214). We are invited to imagine that a social isolate (from birth) finds a Rubik's cube washed up on his island, and learns, on his own, how to solve it. Clearly, solving the Rubik's cube must involve following rules—at least when done repeatedly ('when there is a last corner left to do…'). Ergo, the social isolate can follow rules. The very fact that he solves the puzzle attests to it.

The third example is that of a social isolate who allegedly learns 'the hard way' that he has broken one of his rules:

Robinson junior…is a child who grows up on a desert island. He develops certain methods of coping with his environment. For instance, he makes pots for himself and has found that the juice of a certain berry, when heated, turns red and remains colour-fast. He uses this to 'label' pots in which he keeps dried fruit and he uses other markers to label other pots. Thus, when he 'wants to eat some dried fruit', he finds a red-dyed pot and gets some. (The quote-marks indicate that so far we have only reported his behaviours and not ascribed to him any contentful thoughts, although it aids description to do so.) Imagine further that one day his brain is altered so that he no longer sees red properly (forgive the naïve neuro-psychology and forget 'qualia'). He now has a new disposition to respond and, being a community of one, has no one to check him. But he can still go wrong and discover that he is wrong. He may, even if he does not notice the change, 'think that he has used all his pots of dried fruit' but find that he has not. He may check upon himself by bleeding an animal and find that the relational techniques by which he copes with his domain have significantly changed. It therefore seems possible that he can unilaterally discern a difference between *seems right* and *is right*. (Gillett 1995: 193)

DEFENDING THE STRONG THESIS

I am unconvinced by the individualists' objections. I can be short in my treatment of the first three. Take first the Metaphysical Argument. Clearly the appeal to what is metaphysically possible—what God is able to create—cannot decide the question at issue, since it is precisely the metaphysical possibility over which individualists and communitarians disagree. Individualists believe that socially isolated human rule-followers are metaphysically possible; hence they have no difficulty imagining that God could create a single such creature in isolation. Communitarians regard socially isolated rule-followers as impossible; accordingly they cannot grant that even God could create one. The same must be said about the Internalist Argument. The individualist thinks that at least some normative phenomena are non-relational phenomena of an individual. The communitarian, in contrast, considers all normative phenomena to be relational and social. The Internalist Argument is thus not a real argument: it merely restates the individualist's position. The third argument, the Invention Argument, fails for a different reason. Although it does better than simply restate the individualist's position, it is based on a straightforwardly mistaken assumption. It is simply false that every rule has a first follower, and that this first follower must be the inventor of the rule. Consider the rules of competitive team sports like rugby. Even though the rules of rugby were presumably invented by someone, that someone could not have been following these rules on her own. You cannot follow the rules of rugby on your own, since to follow the rules of rugby is to be following them jointly with others. But once we break the link between inventing and following a rule, the Invention Argument collapses. To invent some rule is not necessarily to be able to follow it on one's own. Only if it were would the individualist have a distinctive argument for his position.

Regarding Parity Argument 1—the parity of group and individual as far as the 'is right/seems right' distinction is concerned—Donna Summerfield has given the appropriate response:

> Some will undoubtedly object that the community, as a whole, will be no more able than an individual in isolation to make a distinction between seeming correctness and actual correctness. After all, there will be no distinction between the way in which we all (or most of us) go on and the way in which we all ought to go on.

> However, there is an important difference between the individual and the community cases. In the case of an individual in isolation, there will be *no* distinction between seeming correctness and actual correctness. In the case of the community, for *every individual,* there will be a distinction between seeming correctness and actual correctness. (Summerfield 1990: 436)

In a footnote Summerfield adds the speculation that Parity Argument 1 is based on 'treating the community as itself an individual, and supposing that there is one way of going on that seems to this superindividual as though it is correct' (1990: 436). I agree with this speculation. The history of theorizing about individual and group is full of misleading assimilations of the one to the other: either the group is conceptualized as an individual writ large, or the individual is thought of as a community writ small (Kusch 1999: 250–1). It is easy to see how both metaphors can mislead. Take the former metaphor first. Once we routinely think of groups as akin to individuals, we will have few qualms about extending psychological vocabulary to groups in a direct and straightforward way: talk of group minds and collective wills will be hard to avoid. This will make it difficult for us to appreciate that all group phenomena are constituted in and through the *interaction* between individuals. Groups consist of interacting individuals. And once we have lost sight of this idea, we have no way of resisting Parity Argument 1.

Assimilating individuals to groups is no less problematic. I submit that this assimilation underlies Parity Argument 2, that is, the claim that, as far as normativity is concerned, *intra*subjectivity is not principally different from *inter*subjectivity. This thought will seem natural and obvious if we think of the individual mind as a group of interacting voices or time-slices. Indeed, many philosophers will find this thought intuitive—after all, it has informed their theorizing as least since Plato's *Republic*. And yet there are reasons for scepticism regarding this assimilation. Surely, a direct negotiation in the present between myself and my family members differs from a negotiation between myself today, myself yesterday, and myself tomorrow. For instance, my family members might sanction (e.g. browbeat) me if I display intransigence, or fail to cooperate. They might persuade me with unexpected arguments, and they might force me to give in. All this is possible because I have only limited control over the course of the interaction and its termination. A negotiation between my different time-slices is a very different kettle of fish. I have no idea how I can browbeat my previous

or future selves, and my present self has full control over the course of the interaction. Does this not show that the two cases are, after all, distinctly different? On this question, I feel compelled to side with Michael Esfeld, who has replied to Blackburn and Gerrans as follows:

> In an interaction of two or more persons, there can be a disagreement in the way in which these persons continue a sequence of examples in a shared environment at a time. That disagreement can trigger a process of negotiating by means of sanctions conditions under which there is convergence. By contrast, in the case of one person who is considered as being in isolation throughout her life, there is no possibility of such a process of negotiating a convergence. The present disposition of the person always has a privileged position; for the past self cannot reply and give feedback. It cannot make available for a person a distinction between what she takes to be correct or incorrect and what is correct or incorrect. The same point applies to the simulation of another person: simulation cannot give feedback that makes available for a person an external perspective in the sense of making available for her a distinction between taking something to be correct and something being correct. (2001: 91)

Finally, we have to address the most important argument (or set of arguments) in the individualist's arsenal, the Similarity Argument. Here too I feel that most of the needed rebuttals have already been put forward. A forceful criticism of the example of Crusoe and the Rubik's cube has been presented by Norman Lillegard (1998). Lillegard rejects Blackburn's suggestion according to which it is 'easy to go through the thought experiment of coming across such an individual' (Blackburn 1984*a*: 84). As Lillegard sees it, this is far from easy. From where, for example, does a congenitally isolated Crusoe get the notion of what counts as a solution, or of anything counting as something?

> Crusoe is sitting on the beach. The cube washes up. He picks it up, fools with it. For years perhaps. Then suddenly it pops into his mind: 'perhaps there is a solution to this'. This is ludicrous. Or, nothing pops into his mind (in particular no English sentences) but he simply starts manipulating the cube, making little marks in the sand, looking back at them now and then, and eventually, he has it! What? The solution! I personally can neither imagine, conceive, nor make sense of this. In particular I do not see how anything that Crusoe did could be considered by him to be mistaken or incorrect, except in the sense of failing to get him from A to B in the way that some other 'move' did. But there could not be anything 'correct' about getting to B itself (though doing so might be pleasing, satisfying, conducive to survival, etc.). We can imagine a born Crusoe finding the solution to the Rubik's

Cube only by forgetting who a born Crusoe is, perhaps, by forgetting that he is a *born* Crusoe. (Lillegard 1998, unpaginated)

The important point is that whatever a congenitally isolated Crusoe does, he is too different from the humans we know for us to be entitled—according to our very own current attribution practices—to regard him as a rule-follower oriented towards independent standards of correctness. We cannot regard him as a rule-follower simply on the ground that he manipulates the cube while 'consulting' his marks in the sand. The marks in the sand might cause him to produce a combination of colours on his cube such that we would count the cube as 'correctly solved'. 'Correctly solved', that is, if *we* had done the solving. Compare the current case with ants that leave odour marks in the forest in order to enable them to revisit a food source. We would not normally regard the ants as rule-followers possessing a sense of correct and incorrect—though they may well be guided by something external to themselves (i.e. the odour marks). We would not regard the ants as rule-followers since we have no idea of how ants could acquire concepts like 'correct' and 'incorrect'. We know how we acquire these concepts, and where they figure: in discursive practices of teaching and training, and in justifications and negotiations with others. Put differently, in order for us to think of someone as a rule-follower, we must be able to think of 'correctness' and 'incorrectness' (or their analogues) as concepts that they possess (as actors' categories) (cf. Bloor 1997: 105). There is nothing in the Crusoe-cum-cube story that invites, or allows for, this hypothesis. In this story we find plenty of regularities, and acting in accordance with one of our rules—but we do not find rule-following.[80]

Similar comments are appropriate concerning the pattern (—…—…) example introduced by Baker and Hacker. Meredith Williams's reply is particularly forceful:

[80] This observation blocks Millikan's (1990) attempt to generate normativity out of biological functions. According to Millikan it is evolutionary history that determines the function of organs and actions, and hence also what we ought to do. As Bloor (1997: 105) rightly observes: 'This line of argument does not really allow us to avoid the appeal to society as the basis for normativity. The reason is very simple. The idea of getting a piece of rule following right, or the idea of getting it wrong in particular cases, are *actors' categories*. They are ideas invoked by rule followers, or would-be rule followers, about one another or themselves. Normative ideas cannot be external to the thoughts and actions of a community of rule followers in the way a fact about its historical evolution might be. The actor's own awareness of these norms is constitutive of their very existence as norms…'.

given the impoverished world that Robinson Crusoe inhabits, there is no way that we can distinguish corrective behavior from a modification of the rules or from the termination of one game for another or from the introduction of an exception permitted by the rule. Perhaps the 'and so on' implicit in the master-pattern warrants systematically altering the pattern inscribed, so that the four dots Crusoe put down initially were the correct continuation, and Crusoe's gesture that we took as a sign of annoyance was rather a sign of rebelling, his examination of the master-pattern and his subsequent behavior a rejection of that master-pattern for another. Other stories could be told that change the interpretation of his behavior and the rule he is purported to follow. That his behavior is seen as corrective depends upon how *we,* as a matter of course, would take the master-pattern.…But what about Robinson Crusoe himself? Would he count his behavior as corrective, that is, as constrained by something other than his own whim or amusement in drawing the pattern of dots and dashes? From his point of view, he engages in a repetitive behavior derived from a series of dots and dashes initially drawn on the wall, 'derived' in the sense that he reproduces the same series of dots and dashes as produced in the original sequence. Nothing more can be said of his behavior and its relations to the so-called 'master-pattern' without reintroducing the notion of his engaging in an act of interpreting the pattern, and that, of course, has already been rejected as an account of rule-following. (Williams 1991: 111)

This leaves us with Gillett's 'Robinson junior', who put his dried fruits into red pots before losing his ability to pick out red from other colours. Allegedly, the isolated Robinson junior is able to arrive at the distinctions 'right versus wrong' and 'is right versus seems right' on the basis of his disappointment and surprise when noticing that all his dried fruits are gone. I fail to see why this disappointment should give rise to an understanding of correctness. Simply the disappointment of not finding the dried fruits where expected is not enough. Nor is it clear to me why Robinson junior should grasp the difference between 'seems right' and 'is right' upon bleeding an animal. Assume that Robinson's changed eyesight makes him see red things as green things, and green things as red things. Gillett seems to think that upon bleeding an animal Robinson reasons as follows: 'Ah, the pots with dried fruits seemed to have disappeared, but now I realize that it is only their apparent colour that has changed. The cause of their appearance changing must lie within me—in my deficient colour vision—for the colour of the animals' blood has always been the same as the colour of pots with dried fruits.' I am unconvinced. Anyone who can reason in such a complicated

fashion must surely already have the concept of correctness; he does not need to acquire the concept on the basis of this reasoning!

STRONG (B AND C) OR STRONGEST (A) COMMUNITY THESIS?

Above I have defended the strong theses of community dependence (A, B, C) against individualistic criticism. It remains for me to explain why I favour A, the strongest thesis, over the weaker ('strong') theses B and C. Most communitarians regarding rule-following are advocates of C. This is clear from the fact that almost always it is the *congenitally* isolated individual that is the central focus of discussion. Advocates of the Past-Tense Community Thesis C are relying on the following thought. A congenitally isolated individual cannot acquire the conceptual distinction between 'seems right' and 'is right'. However, once an individual has learned to make this distinction in and through interaction with others, the distinction remains permanently available to the individual—even if he finds himself in long periods of physical isolation, or even if he decides to introduce a new rule for himself. Advocates of the Strong Present-Tense Community Thesis B are proposing a synchronic version of the same general idea. If only the distinction between 'seems right' and 'is right' is generally available in a community, and if an individual has learned the distinction from others, then the individual can apply it to his own rule-following, never mind whether the rule is shared with others in the same community.

I disagree with these lines of thinking. When we ask whether an individual A can follow rule R *now,* we are asking whether A can *now* meaningfully distinguish between 'seems right' and 'is right' with respect to actions governed by R. And for this it is not enough to point to A's ability to distinguish 'seems right' from 'is right' with respect to some other rule, or to A's *past* ability to distinguish 'seems right' from 'is right' with respect to R. Consider again the distinction between rule-stipulation and rule-application. I suggested earlier that this distinction is not clear-cut. Paradoxically put, in applying a rule, we stipulate its content. In deciding how to apply my rule 'go running in the morning' I decide which rule it is I am actually following. Clearly, as long as I do so on my own, I am not encountering any friction; however I choose to apply it my rule is correct, since

it is only my decision regarding application that fixes (momentarily) the rule's content. In other words, whatever seems right to me is right. And this means that the distinction between 'seems right' and 'is right' does not have a foothold in my practice. It can apply to me only if a gap can open up between what I decide to allow as correct and what others are willing to accept as correct *with respect to my rule of going running*. The problem of generating a distinction between 'seems right' and 'is right' for this rule is not solved by having the distinction available for some other rule. My family might have a rule to go out for dinner on Sundays, and, given that we interact and correct each other in applying this dinner-rule, I have access to the distinction between 'seems right' and 'is right' for this rule. But again, having the distinction available here does not make it accessible to me if then I try to follow the running-rule on my own.

But can I not solve the Rubik's cube on my own today—even though everyone else has long since moved on to more exciting computer and board games? Can I not distinguish between what seems like the correct solution and what is the correct solution? Maybe the cube rests on the table, and maybe the three visible faces are each of a single colour. But when I turn the cube round, I realize that some faces of the cube are not of one colour. What *seemed* like the right solution *is not* the right solution. Moreover, I certainly know what solving the Rubik's cube consists in. Here then, it seems, there is a clear-cut separation of rule-stipulation and rule-application even for the solitary rule-follower—provided only he is in general able to understand the distinction between 'seems right' and 'is right'.

Closer inspection shows that things are not as they seem. In fact, the case of the Rubik's cube does not differ at all from the case of the morning run. To tell myself 'try to end up with a single colour on each face' is like telling myself 'go running in the morning'. The content of the latter rule is determined (temporarily) only as I make further decisions about how to follow the rule in different circumstances. In the case of the Rubik's cube these include the following: Do I wish to end up with a single colour on each face of the Rubik's cube? Or do I wish to end up with a particular pattern of different colours on one or several of the faces? Do I have to achieve a given pattern in a particular number of moves? Less than a hundred? No more than sixty-five? Do I have to achieve a particular distribution of colours within a specific time period? Am I allowed to consult literature with advice