

For Khani and Kemp, eds, *Wittgenstein and Other Philosophers*, Vol 2

Pettit's Wittgenstein

Philip Pettit

It feels like hubris to write about my views and Wittgenstein's, as if they were of equal interest and importance. But signed up as I am, I have no choice but to satisfy the commitment I made as this volume was being planned. I trust that the gods will not take umbrage.

1. Discovering Wittgenstein

I came to Wittgenstein from a rather unusual background. Trained primarily in the continental rather than the analytic tradition, as it was common to describe those rival approaches in the 1960's and 1970's, I was devoted initially to the phenomenological way of thinking represented by Edmund Husserl, although I was more enthused by the work of those who followed him like Sartre and Merleau-Ponty and, to a lesser extent, Martin Heidegger. The catch cry of phenomenology came to worry me, however, and eventually led to a loss of faith in the approach. This was Husserl's call, *Zu den Sachen selbst*: go back to the world of experience, putting aside all the presuppositions of theory and tradition.

One of the works that gave me second thoughts about the direct appeal to experience was an early book by Jacques Derrida, now not often cited by his followers: *La Voix et le Phénomène*. In a rather more moderate version of the views that later made Derrida a celebrity, he argued that Husserl neglected the extent to which anything given in experience was bound to be structured by the language of the investigator and by the assumptions it encoded in its distinctions and connections. This rang true with me, echoing then familiar themes in the philosophy of science according to which observation was often theory dependent. Equipped with suitable concepts, to take a common example, the physicist could see, and perhaps be bound to see, electrons in the traces detectable in a Wilson cloud chamber.

I quickly went off Derrida, and the structuralist approach he initially seemed to exemplify, because of his obsession, as it seemed to me, with wordplay: often amusing and insightful, it just didn't seem serious. My retreat led, I cannot recollect why, to a reading of

the later Wittgenstein. I had known the *Tractatus* fairly well, having done an MA dissertation on the early work of Moore and Russell with which it was connected. But reading the *Philosophical Investigations* and his other later writings opened wholly new vistas for me.

What the *Investigations* made clear to me was that the human mind did not disclose itself in the sort of epiphany that Husserl, or at least Husserlians, had come to seek. Our success or failure to understand how our minds work, I came to think, depended on examining the sorts of words we brought to the articulation of mental phenomena, such as 'state' or 'process', 'thought' or 'perception' or 'insight'. And it depended equally on how far we were captivated by pictures that suggest themselves naturally, at least when we begin to think in a philosophical manner. Initially I was even taken by Wittgenstein's idea that a rigorous philosophical discipline could provide a therapy by means of which we might hope to escape capture by such images.

I quickly lost this belief in philosophy as therapy and came to think that we could learn from Wittgenstein's many insights, while seeking to build them into something more like a traditional philosophical theory. This would be a view of things that started from certain assumptions or axioms and evolved on the usual argumentative lines into a general account of the topic addressed. In breaking in this way with the therapy that Wittgenstein recommended, I felt that I could still retain many of his crucial ideas.

There are two claims that I took away from my reading of Wittgenstein and I seek to address them in sections 2 and 3. The first is that there is a deep puzzle at the very basis of human thinking—the rule-following problem, as it has long been known—that had not been properly addressed in philosophy before Wittgenstein took it up. And the second is that the resolution of that problem lies in recognizing that we human beings can follow rules of the kind presupposed by the exercise of thought and reasoning only in virtue of being engaged in social practices with one another.

This second theme excited me greatly, perhaps because it answered to a personal intuition that we are social animals in a deeper sense than even Aristotle recognized in describing the human being as a *zoon politikon*. I have always been surprised at those who

recoil from that image, apparently preferring to think of human beings in Cartesian terms as mental loners: creatures who reach out to one another's minds from positions of an independently accomplished competence in thinking and reasoning, albeit a competence causally elicited in childhood. Perhaps because of coming from a densely social background in rural Ireland, the social image of mind that appears in Wittgenstein appealed to my temperament.

I worked on the Wittgensteinian material in the 1980s and 1990s, inspired though not fully persuaded by the insightful commentaries from Crispin Wright (1980) and Saul Kripke (1982), and I published a range of this work in the following years (Pettit 1993; 2002, Pt 1). Having concentrated in the meantime on issues to do with the nature of freedom, group agency, value and responsibility, I returned to thinking about the social nature of mind in the last five years or so and am currently writing up a book based on the Locke lectures that I gave in Oxford in 2019 (Pettit 2024).

In this recent work I rely on a method that may itself have echoes in Wittgenstein's writings. He often asks us to imagine creatures with dispositions and practices very different from ours, using that thought experiment to highlight how differently they would think of things and to suggest that our ways of thinking depend in the same way on our own dispositions and practices. In my book, employing what is now often called a pragmatic genealogy (Queloz 2021), I argue that creatures otherwise like us would be robustly likely to respond to the advent of a basic language, and to living under linguistically enabled practices, by developing novel, distinctively human habits of thought and reasoning. And on that basis, I try to defend a theory of the human mind that casts many of our distinctively mental capacities as the interiorization of primarily social skills. But time to return now to the two themes that I promised to address.

2. The rule-following problem

Following rules, basic and otherwise

Let me outline what I see as the rule-following problem before turning to Wittgenstein's own presentation of it.

A distinctive feature of our species is that we ask one another questions, even ask ourselves questions, and try when the evidence is available to generate answers. Among the questions raised, we ask about whether properties (and related entities) that we cannot analyze or define—for short, basic properties—are realized in this case or not. And without having personal access to definitions, we try to provide answers. Is that a *tool* in your hands? Is that a *game* that they are playing? Is this number the *sum* of those? Is that coat *red* in color? Is that a *regular* shape or not? That we lack personal access to definitions in such cases does not mean that the properties are indefinable; they may be basic-for-us without being basic-for-experts. But some properties—some suitably determinate properties—must be basic for us, on pain of regress, as indeed some must be basic for experts. And we may assume that the cases given illustrate that category.¹

When we try to answer a question of this kind, the manifest assumption is that we understand the property involved, being aware of the sorts of conditions determining whether it is present, and that the aim of the exercise is to check on whether they are realized according to the evidence at our disposal. We think of that exercise as one in which we may fail but, at the same time, as one in which a greater effort on our part—paying closer attention to the property and to the evidence at hand—can reduce the likelihood of failure.

That there are stable conditions linked with the property, determining whether it is present or not, means that there is a regularity in that linkage. That regularity will serve intuitively as a rule governing judgments about its presence, insofar as we are able consciously to try to conform to it—equivalently, able to try to track the property—expecting that the attempt may raise our chances of success but not necessarily ensure success. We cannot try to conform to it by consulting a formula that lays out the realization conditions of the property, since such a formula would require the sort of analysis or

¹ It may be, for all we assume, that what is basic-for-us or basic-for-experts are inter-defined packages of properties, not properties in isolation: this, in the way a line is defined, roughly, as the shortest distance between two points and a point as the intersection of two lines. We ignore that possibility here for reasons of convenience.

definition that we are taking to be absent. We can only try to conform to it by seeking to track the property directly, looking for evidence of its realization in this or that situation.

Modes of rule-following

When a regularity serves as a rule in this sense, then we control for conforming to it in a conscious and intentional manner; we manifestly try to ensure conformity. That the control is conscious follows from the part that the property plays as an object we identify as a target to track; that it is intentional is implied by the need for effort in this tracking. The fact that such control is necessary to promote conformity to the rule, and that control is never perfect, means that besides being conscious and intentional, rule-following has the further feature of being defeasible, indeed defeasible in a way that must be salient to the agent. Despite our best efforts, we may fail to conform: the control may be unsuccessful.

Is this conscious and intentional mode of control common? No, but it is nonetheless fundamental.

It is not common because we often control in a standby rather than an active fashion for following a basic rule. In those cases, we let unthinking habits dictate the judgments we make, and the words we utter, in tracking the instantiation of this or that property. But we retain standby control over the operation of those habits insofar as two conditions hold. First, a red flag will go up, attracting our attention, if the habits prompt an implausible judgment: say, a judgment that something is a tomato when it turns out to taste spicy. And second, that red flag will prompt us to resort to conscious, intentional effort in determining what it is that we put in our mouth; it may lead us to conclude, for example, that it is a pepper. In such a case we retain control to the extent that we are disposed to notice a red flag—‘Is this really a tomato?’—and to let that flag prompt us to think reflectively about the case.

It is likely that on most of the occasions when we follow a basic rule, we do so without thinking, as we naturally say: we let habit take over. But if the intentional, conscious mode of rule-following is not common, as this implies, still it is fundamental. This is because standby rule-following is possible only if active rule-following is possible: only if

it is in the wings, waiting to be activated. For that reason, it is the active mode on which we focus here, as it is the active mode that engages Wittgenstein too.

The rule-following problem formulated

How do we manage to directly track a property that is basic for us in the sense explained: a property or indeed any similar entity, such as the plus-function implied in the property of being a sum?² In the definitional case, the resources that enable us to control appropriately for conformity are provided by the formula available. The problem in the case of following basic rules, specifically the active version of this case, is to explain what the resources are that enable us to track a property directly rather than definitionally.

The idea that certain rules are basic connects with the idea of basic acts. That an act is basic means that while we perform it intentionally, we do not perform it by performing any other acts intentionally (Hornsby 1980). I may fasten my shoes intentionally by tying my shoelaces intentionally. But if I am proficient in doing so, I will tie my shoelaces intentionally without relying on doing anything else intentionally. The child may have to learn to move its fingers, now in this way, now in that, to tie its shoelaces, so that the act of tying is not basic for it. But when the child becomes proficient, it will tie them intentionally without any awareness of what it does with its fingers, and so without intentionally moving those fingers in any independently characterized way. It will tie its shoelaces intentionally; and it will tie them without relying on intentionally taking a distinct step as a means to that end.

Actively following basic rules is not only an intentional act but a basic intentional act. Following a rule that is defined for us in other terms means intentionally conforming to the rule in suitable situations by means of intentionally applying the relevant definition or formula. Following a basic rule means intentionally conforming to it in response to appropriate evidence but not by means of doing anything else intentionally—anything more basic—such as applying a definition.

² Rule-following in familiar cases involves explicit or implicit definitions or formulae and, while it too may raise some problems, we shall concentrate here, following Wittgenstein, on the basic case where definitions are lacking. Paul Boghossian (2012) focuses on a more general problem, which we shall ignore here. For a response, see (Pettit 2024, Ch 2).

With plausible examples of basic rule-following, such as those given earlier, it may be hard to imagine that any creature could try to track a property without having a word to refer to that property or, equivalently, a word to express the concept under which it is viewed. Trying to track the property may consist as a matter of practice in trying to use a relevant word only in the presence of the appropriate property. But even if the effort to track a property is not necessarily connected with having a word for it, still it is worth noting that if an agent could use a word in a way that was guided by an awareness of the corresponding basic property, that would mean that it was tracking the property in rule-following sense, judging that it to be present in certain cases and absent in others.

The problem in Wittgenstein

The rule-following issue, as we have described it here, is at the center of Wittgenstein's discussion, and various commentaries on that discussion. Kripke (1982, 24) puts the problem nicely when he formulates the requirement for the direct tracking of a property or other basic entity. The requirement is normative, and manifestly normative, in character: viz., that the property involved 'should *tell* me what I ought to do in each new instance'. If it didn't speak to me in some such sense, he suggests, then to judge that the property is present in a this or that case would be to 'make an unjustified leap in the dark' (Kripke 1982, 10).

Wittgenstein's (1958, para 175) discussion highlights the challenge of explaining basic rule-following in much the same way, arguing that the target involved—the property tracked—should *guide* me, as he puts it. It may seem that in following a rule I was just moved to go one way rather than another. But that seems wrong, he writes: 'I feel as if there must have been something else'. "For surely," I tell myself, "I was being *guided*". The problem, then, is to say how a basic property could *guide* me across cases: how, in Kripke's phrase, it could *tell* me which cases are instances of the property, which not.

This distinguishes the rule-following problem from something that Jerry Fodor (1987; 1990) called the disjunction problem. Suppose we take a human agent to represent some particular as an instance of a general property: say, that of being a game or a tool. Grant that any such agent may make a mistake, on our ordinary view of representation:

say, that they may take a fight to be a game or a mere ornament to be a tool. On what basis do we judge that the agent misrepresent things in making such a judgment? Why do we not exercise charity and assume that the content of the judgment or representation—the meaning of the word employed, if there is a word—is that something is a game-or-a-fight, a tool-or-an-ornament?

This problem requires us to resolve a certain normative issue: Why is it correct to take the representation to have the non-disjunctive content? But this problem is not one of normative guidance. It is a problem about how right and wrong gets to determined—a problem that arises for any simple animal or robot to which we ascribe representations—not the more complex problem of how right and wrong get to be determined in such a way that agent can be guided to satisfy the desire to do what is right.

It would be unilluminating, according to Wittgenstein, to say in the sort of example he has in mind that the agent just grasps the nature of the property addressed, the sense of the concept or word that ascribes it (see too Kripke 1982, 54). That may seem to be an attractive gloss, as he admits: "It's as if we could grasp the whole use of a word in a flash" (Wittgenstein 1958, para 197). But how is that grasp, that intuition, to guide me, if I am that agent? Perhaps by an 'inner voice', he wonders? But the suggestion that there could be such a voice comes to naught: 'how do I know *how* I am to obey it? And how do I know that it doesn't mislead me? For if it can guide me right, it can also guide me wrong' (Wittgenstein 1958, para 213).

Putting intuition aside, Wittgenstein asks whether I might grasp a basic property by surveying a set or series of instances and then extrapolating from those to other instances. Might I cotton on to the property by looking at instances, for example, of 'the same colours, the same lengths, the same shapes' and thereby learning to 'continue' the 'pattern uniformly' (Wittgenstein 1958, para 208)? No, he claims. For how am I to know how 'to continue a pattern?' (Wittgenstein 1958, para 211); 'how can a rule shew me what I have to do at *this* point?' (Wittgenstein 1958, para 198). Might I find reasons to go this way rather than that? No, for 'my reasons will soon give out' and then I can only 'act, without reasons'.

The problem is that there is nothing about a finite series of any items that gives me reason to think of extrapolating to further items as following a rule. Any finite set of would-be examples of a rule—say, examples of games or tools—will fail to determine whether some new item is or is not another example. ‘Whatever I do is, on some interpretation, in accord with the rule’ (Wittgenstein 1958, para 198).

It is plausible with any series of items, of course—say, examples of addition which illustrate one or another number as the sum of others—that we develop a disposition to continue in one way rather than others; that is part of what happens when we learn to add. So perhaps rule-following just consists in forming such a disposition and then acting as it prompts me to act when I consider further instances? Perhaps following the rule for detecting sums, for example, just involves being ‘disposed, when asked for any sum “ $x+y$ ” to give the sum of x and y as the answer’ (Kripke 1982, 23). In a plausible interpretation of Wittgenstein, Kripke (1982, 24) argues that this won’t work either, principally for the reason that to be subject to a disposition in proceeding is not to be guided towards what one ought to do or is justified in doing; it conflicts with the assumption in such a case that ‘whatever in fact I (am disposed to) do, there is a unique thing that I should do’.

At this point we are at an impasse. Any basic rule is meant to cover an indefinite range of cases, and to guide that agent in each case to identify what it is right to think or say there. The rule is meant to operate in the way we might imagine a defined rule operating, say the rule defining multiplication, in guiding us on what is a case of multiplication and what not. Let multiplying a number m by another number n be defined as adding m to itself n number of times. We might follow that defined rule by checking in each instance whether the number in question relates in the appropriate way to any other two numbers. But consulting that defined rule involves relying on other more basic rules, such as that associated with addition. And the rule-following problem is that there is nothing like a definition—nothing like an interpretation of the rule—that is available here for consultation.

The upshot is clear. If we assume that active rule-following involves something like consulting a definition or interpretation as we look for the right response in each new instance, then we just cannot make sense of basic rule-following. The lesson, in

Wittgenstein's (1958, para 201) words, is that 'there is a way of grasping a rule which is *not* an *interpretation*, but which is exhibited in what we call "obeying the rule" and "going against it" in actual cases'.³ But the question in that case is what it might be to follow a rule if it is not to be guided by an interpretation of the rule?

3. *The social resolution of the problem*

Wittgenstein responds to the problem in a somewhat aphoristic and opaque manner. He suggests that the rule-following problem leaves us in the lurch only because we are thinking about it in the wrong way: 'we are like savages, primitive people, who hear the expressions of civilized men, put a false interpretation on them, and then draw the queerest conclusions from it' (Wittgenstein 1958, para 294). But unfortunately, he is very elusive about what exactly the right way of thinking involves.

Thus, he compares a rule to a sign-post, suggesting that the solution is to see rule-following as governed by social convention: 'a person goes by a sign-post only in so far as there exists a regular use of sign-posts, a custom' (Wittgenstein 1958, no 198). Emphasizing the social aspect, he adds that "'obeying a rule" is a practice', so that 'it is not possible to obey a rule "privately"' (Wittgenstein 1958, 202). But how exactly to understand the approach that these remarks, and many similar comments, suggest?

The comments that follow represent my attempt to make sense of basic rule-following, but I do not offer them as a scholar of Wittgenstein's work. In its earlier formulations, I frequently presented my approach in that way, only to find that I elicited strong resistance, including remarks to the effect that Wittgenstein was much more nuanced and sophisticated than I took him to be. When I stopped presenting the approach in that manner, however, describing it as my own attempt to grapple with the issues, I continued to encounter resistance, this time in the form: 'But that's all in Wittgenstein!'. Chastened by these responses, I can only present the paragraphs that follow as a view that I take to be grounded in Wittgenstein's insights but that I do not dare to offer as an informed interpretation.

³ For a fine account of Wittgenstein's rejection of the role of interpretation in basic rule-following, see (Miller 2015). See also (Swindlehurst 2020).

There are three practices that are common to human beings, both as a matter of common sense and of science, and I think that we manage to follow basic rules in virtue of being immersed without thinking in those practices, and in virtue of being committed to assumptions that they encode. I describe the three practices as those of sensitization, joint action and teaching-learning.

Sensitization

Take any perceptual similarity class like that of red colors or flat surfaces or pear shapes and assume that the properties involved are basic for relevant agents. Common sense and laboratory experiments show that even simple animals are capable of recognizing the instances of such a similarity class, as when they are conditioned to respond to an instance: this, for example, in the way a pigeon can be conditioned to peck at pear-shaped boxes, being rewarded in each case by an tasty seed. In such a case we can say that the agent is sensitized to the class of things in question.

When a creature as simple as a pigeon is sensitized to pear shaped boxes, then it forms a belief in any instance that the box is of an appropriate kind and is led by that belief to peck at the box in order to satisfy its desire for food. This process of belief formation presumably involves a psychological attunement that is not available to the animal's consciousness or within its control. Still, it means that there is a sense in which the pigeon believes that in suitable contexts all pear shapes are ways to food but it believes this *in sensu diviso*, as it used to be said, not *in sensu composito*. Such a belief amounts to nothing more than a disposition to believe of each pear shape that it offers a way to food and does not require the pigeon to have an idea or concept or representation of the pear shape as a general kind, let alone a notion of universal quantification. It will be aware as such of each particular shape presented, being able to form an action-selecting belief about it. But the pigeon need not be aware of the kind as such: this need not be an object it can consider, an entity about which it can form any beliefs..

Does sensitization of the sort exemplified in the pigeon's behavior enable it to follow the basic rule associated with being pear shaped? Clearly not, for while the disposition prompts suitable responses it does not, for reasons emphasized by Kripke, provide

anything that might pass as normative guidance. It is worth noting that the pigeon might even be prompted to use a sound—a candidate for a sign—in response to every pear shape it confronts—or, more plausibly, in response to every instance of property important for pigeons generally—without following a rule. Indeed, it might even be the case, despite the absence of rule-following, that the sound it makes plays the role of a sign amongst other pigeons—say, a sign of a predator—prompting them to respond as if they had perceived what the signing pigeon perceived.⁴

Human beings, like other animal species, are certainly capable of being sensitized to classes that will present to them as similarity classes. But equally certainly this in itself will not involve them in the following of basic rules. There is a second practice that goes beyond triangulation, however, and promises to do better in crossing the gap between brute disposition and rule-following. This is the practice of joint action that human beings naturally display, given the evidence from young children and may be restricted to members of our species.

Joint action

Human beings, and hence the humanoids of our narrative, are creatures who spontaneously act together for various common goals, going beyond the sort of action that mere sensitization would support. They have a distinctive capacity and inclination to combine their efforts to advance any goal where it is manifest, first, that they each desire its realization and, second, that they can only or best achieve this in tandem, with each playing their part in a salient plan. Thus, if they are on the beach and they observe that a swimmer is in difficulty, they will be likely to get together to save the swimmer when it is manifest that this is a goal they share, that there is a salient plan under which they can realize it together, and that anyone who begins to enact the plan will be joined by others. They may save the swimmer under such conditions, for example, by getting together to form a chain

⁴ For an example of such signing amongst animals, see the research on vervet monkeys in Kenya that is documented in (Cheney and Seyfarth 1985)

of people into the water; this may be the salient thing to do, perhaps because someone suggests it.⁵

Michael Tomasello (2016) argues that this predisposition towards jointly intentional action is one of the most distinctive features of human beings. Tomasello relies on two sources of evidence to support the claim that joint action is a characteristic of human beings. The first is that in a crucial period of human evolution, between about 400K and 150K years ago, the environment was such that our human forebears would have been forced to forage and hunt together—this, or die alone—and that that would have created a selectional pressure in favor of a natural disposition to act jointly. They would have had to be able to distinguish edible from poisonous plants, and to collaborate in picking the edible and avoiding the poisonous. And equally they would have had to be able to recognize potential prey and potential predators and to combine in hunting the animals of the one sort and in defending against animals of the other.

The second source of evidence on which Tomasello draws is that the disposition to act jointly with others is displayed by children between the ages of 1 and 3, although it is generally not displayed by other primates. ‘These young children coordinate on a joint goal’, Tomasello (2014, 41) says, ‘commit themselves to that joint goal until all get their reward, expect others to be similarly committed to the joint goal, divide the common spoils of a collaboration equally, take leave when breaking a commitment, understand their own and the partner’s role in the joint activity, and even help the partner in her role when necessary’. As examples of such collaboration among very young children he mentions ‘giving and taking objects, rolling a ball back and forth, building a block tower together, putting away toys together, and “reading” books together’ (Tomasello 2014, 44).⁶

In order for any agents to practice jointly intentional activity of this kind, they must not only be sensitized to this or that instance of a class of activity or a kind of object. They must also be able to direct their attention to the class or kind itself. They will have to do

⁵ This is a very stark statement of what is involved in joint action, broadly in line with (Bratman 2014); see (Pettit 2017; 2023, Ch 4). For other approaches to the analysis of this notion, any one of which would work for our purposes here, see (Tuomela 2007; Searle 2010; Gilbert 2015).

⁶ See too (Tomasello 2008; 2009; 2016).

this when they form a belief that a partner is seeking a joint action in a certain class: say, that of playing some sort of game. And they will have to do it when they seek with others to find an object in a certain class: say, a plant of such and such a kind. They must be attuned to the property that unites instances of that activity and instances of that object. And, more than that, they must assume that their partners in the enterprise are also attuned to the property and that the partners assume the same thing about them.⁷

If agents did not have this extra capacity, then they could hardly plan to pursue a certain class of animals with others or search out a certain class of plants. Indeed, they could not plan to engage with others in any class of activity, even something as simple as playing a game together. Thus, without that capacity, to return to Tomasello's case, no child could expect collaboration on the part of another in rolling a ball or building a tower or reading a book. And no child would have grounds to remonstrate with another, as is apparently common among children, about their not enacting the pattern properly: not adding to the tower, not rolling back the ball, not joining in looking at a book, and so on.

We saw that sensitization to a class or property or any such abstract entity will not enable agents to have the general kind as an object of consideration about which they might form general beliefs *in sensu composito*. This all changes, we can now see, with the advent of joint action, for it will require them, not just to be responsive to instances of a kind, but to have a notion of the general kind as such. They will have to be able to form beliefs to the effect that this activity is an instance of the same kind as some previous instance, or that this is an object of the same kind as on which they acted before. And, moreover, they will have to be able to assign similar beliefs to others.

The difference made by the capacity for joint action shows up in the signs that we can imagine them coming to use in response to instances of one or another kind about which they can form beliefs. Suppose that they come to use 'tigroo' of animals in a certain class; we may suppose that this comes to them, not by natural instinct, but by social

⁷ This consideration may also support the stronger claim that the assumption has to be manifest or a matter that is public between them. Such manifestness may be interpreted for current purposes as involving common awareness in the sense of David Lewis (1969): that is, a hierarchy of assumptions involving the assumption by each, not only that they all assume a commonality, but that they all assume that they all assume the commonality, and so on. On the case for rival interpretations see (Lederman 2018).

contagion and convention. Since they will be attuned to that class as such, and the property common to members of the class, the sound ‘tigroo’ will presumably function for them as a marker of the property. And that means in turn that an utterance like ‘tigroo’ amongst them will naturally be taken by them and by others—assuming no incompetence or insincerity—as an expression of the belief that there is a tigroo present: this belief that may even be important for them in the organization of joint action.

Would the use of such a sign testify to their following a basic, common rule? Not perhaps if it were wholly non-intentional: not if it occurred as a matter of instinctive response. But we may presume that it can be brought under intentional control, as even a natural as distinct conventional sign might be intentionally controlled. Would it in that case demonstrate that in using the sign, those involved in the practice must each be intentionally tracking instances of the kind to which it presumptively applies?

Not necessarily, as appears when we ask about how two human beings in our scenario might be expected to respond to a discrepancy in their respective usage of such a sign, and presumptively a discrepancy between the beliefs they each form in that situation. There is no reason in the story so far why either of them might balk at the divergence, as they would presumably do if they could be cast as intentionally and consciously seeking to follow a common rule. For all that we have assumed, they may simply turn away from such a conflict in their signaling and in their beliefs; they may just give up on the joint activity that convergence would likely have triggered.

The idea is that for all we have assumed, the participants in the signing activity may each take the property picked out in their signing, ascribed in their belief, as one that is canonically identified in their personal disposition to hold, now in one case, now in another, that such and such an animal is or is not tigroo. The activity may not have the fallible, let alone manifestly fallible, character of an exercise in rule-following. They will not be trying to follow a rule that they might miss or mistake; they will be playing their own spot-the-tigroo game, not a common game, and it will be a game in each case that they cannot lose.

Teaching-and-learning

Class to guide fallibly

This being so, the activity in which they are each involved falls short of rule-following proper.

References

- Boghossian, P. (2012). "What is Inference?" Philosophical Studies **169**: 1-18.
- Bratman, M. (2014). Shared Agency: A Planning Theory of Acting Together. Oxford, Oxford University Press.
- Cheney, D. L. and R. M. Seyfarth (1985). "Vervet Monkey Alarm Calls: Manipulation through Shared Information." Behaviour **94**: 150-66.
- Fodor, J. (1987). Psychosemantics. Cambridge, Mass., MIT Press.
- Fodor, J. (1990). A Theory of Content. Cambridge, Mass., MIT Press.
- Gilbert, M. (2015). Joint Commitment: How we Make the Social World. Oxford, Oxford University Press.
- Hornsby, J. (1980). Actions. London, Routledge.
- Kripke, S. A. (1982). Wittgenstein on Rules and Private Language. Oxford, Blackwell.
- Lederman, H. (2018). Common Knowledge. Handbook of Social Intentionality. M. Jankovic and K. Ludwig. London, Routledge.
- Lewis, D. (1969). Convention. Cambridge, Mass., Harvard University Press.
- Miller, A. (2015). "Blind Rule-Following and the "Antinomy of Pure Reason"." Philosophical Quarterly **65**: 396-416.
- Pettit, P. (1993). The Common Mind: An Essay on Psychology, Society and Politics, paperback edition 1996., New York, Oxford University Press.
- Pettit, P. (2002). Rules, Reasons, and Norms: Selected Essays. Oxford, Oxford University Press.
- Pettit, P. (2017). Corporate Agency—The Lesson of the Discursive Dilemma. Routledge Companion to Collective Intentionality. M. Jankovic and K. Ludwig. London, Routledge.
- Pettit, P. (2023). The State. Princeton, NJ, Princeton University Press.
- Pettit, P. (2024). When Minds Speak: A Social Genealogy of the Human Soul. Oxford, Oxford University Press.
- Queloz, M. (2021). The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering. Oxford, Oxford University Press.
- Searle, J. (2010). Making the Social World: The Structure of Human Civilization. Oxford, Oxford University Press.
- Swindlehurst, Z. M. (2020). "Blind rule-following and the regress of motivations." Inquiry **64**.
- Tomasello, M. (2008). Origins of Human Communication. Cambridge, Mass, MIT Press.
- Tomasello, M. (2009). Why We Cooperate. Cambridge, Mass, MIT Press.
- Tomasello, M. (2014). A Natural History of Human Thinking. Cambridge, Mass, Harvard University Press.
- Tomasello, M. (2016). A Natural History of Human Morality. Cambridge, Mass, Harvard University Press.
- Tuomela, R. (2007). The Philosophy of Sociality: The Shared Point of View. Oxford, Oxford University Press.
- Wittgenstein, L. (1958). Philosophical Investigations. Oxford, Blackwell.
- Wright, C. (1980). Wittgenstein on the Foundations of Mathematics. London, Duckworth.