Article:
Dodd, J orcid.org/0000-0003-1737-7616 (2023) 'Rails Invisibly Laid to Infinity'. The Philosophical Quarterly, 73 (1). pp. 84-104. ISSN 0031-8094

https://doi.org/10.1093/pq/pqac012

21.03.2022

‘RAILS INVISIBLY LAID TO INFINITY’

BY JULIAN DODD

This paper addresses what I call ‘the constitutive question’ concerning the rules we follow: namely, what determines the standard for a rule’s correct application. John McDowell has offered a putative ‘middle position’ between two extreme, unacceptable answers: *empirical idealism*, which takes the requirements of a rule in any given situation to be constituted by our reaction to the case; and *hard platonism*, which takes these requirements to be delivered by unvarnished reality as absolutely the simplest or most natural way to carry on. Tellingly, however, McDowell’s position is itself unacceptably idealist in his picture of the way in which we are ‘involved’ on the right-hand side of biconditionals such as ‘“Diamonds are hard” is true if and only if diamonds are hard’.

In response to this stultifying state of affairs, I suggest that McDowell has followed the empirical idealist and hard platonist in assuming that the requirements of a rule must be grounded by *something else*: something external to the rule, which mediates between the rule and the standard for its correct application. This assumption is false. What is more, one application of a metaphor to which Wittgenstein has a somewhat equivocal relation – that of rules as rails invisibly laid to infinity – can help this point to stick.

Keywords: rules, rule-following, Wittgenstein, idealism, platonism, McDowell

1. ‘Normative reach’, objectivity, and the pull of idealism

I start with a familiar observation. Understanding words is a matter of having grasped the rules for their correct use. Someone who understands the instruction to keep adding 2 has grasped the rule its words express for developing a numerical sequence; someone who understands an everyday descriptive word, like 'hard', has grasped the rule for applying that word in accordance with the concept it expresses. She knows what she has to do in order to follow it.

Now for something that is less of a commonplace. We should distinguish three questions concerning the rules we follow. Here are two of them. First, there is the *epistemological* question: namely, how is it possible for us to follow such rules? How does a rule manage to lead us (Wright 2007: 482)? Second, albeit relatedly, there is the *relevance* question. This is the question scintillatingly pressed by Saul Kripke (1982), and which has dominated discussion of rules and rule-following for forty years or so: what, if anything, determines that we are following one particular rule, when everything that we have thought, said or done up to now is consistent with our following any of an indefinite number of potentially extensionally divergent rules (Wright 2007: 482)? How can it be *that* rule that we are led by?

Wittgenstein offers us a picture of rule-following – summed up by the observation that following a rule is a *custom* or *practice* (Wittgenstein 1953: §§199, 202) – that is designed to take the sting out of these two questions. We

are led by *that* rule, and not by one of its potentially extensionally divergent competitors, because, given both our shared nature and common training in the use of words, it is this rule's pattern of correct use that counts for us as 'carrying on in the same way as before'. Being the people we are, embedded in a particular system of practices, the potentially extensionally divergent rules do not present themselves as options for us. And it is precisely because this is so that the epistemological question can be tamed. Nothing special of us is required for us to follow a rule: we just follow rules by using words in the ways in which have been trained. In particular, we can follow a rule without representing to ourselves how to do so: there is a way of grasping a rule that is '*not* an interpretation, but which, from case to case of application, is exhibited in what we call "following the rule" and "going against it"' (1953: §201). This achievement is aptly described as being able to follow a rule 'blindly' (1953: §219).

But none of this speaks directly to the question that I shall focus on. Following Crispin Wright, I call this third question *constitutive* (Wright 1989: 305; 2007: 488). The constitutive question concerns not how we can be led by a rule, or how it can be settled which rule we are following, but what constitutes – we might say 'determines' – the standard for a rule's correct application (1989: 304; 2007: 488). Imagine, then, that someone is instructed to develop a numerical series by repeatedly adding 2, and that having grasped

the rule for doing so, she expands the series in the expected way until she reaches 1000. What *makes it the case* that the rule she is led by requires her to continue '1000, 1002, 1004, 1008 ...', and not '1000, 1004, 1008, 1012 …' (Wittgenstein 1953: §§185-187)? Figuratively, we can think of the rule being followed here as something which, step by step, guides its followers to develop the series in the former direction, rather than the latter (Wright 2007: 497). But what determines that *this is* the grasped rule's direction, that *this is* the way in which its pattern of correct application extends? Do facts about ourselves and our linguistic practices figure in an explanation of what the rule's requirements are, or are these requirements independent of such facts?

The constitutive question can be hard to keep clearly in view. Its subject matter is not how our uses of words determine *which* rules these words express. That is Kripke's relevance question again.[1] The constitutive question, by contrast, operates at the level of the rules themselves: at the level of our concepts. It asks of *any* concept that we might use a word to deploy, how *its* standard of correct application is determined (McDowell 1984: 256, n. 4; Wright 2007: 487-488). And it is precisely because it is a question operating at

[1] There are terminological elephant traps here. Wright, as I have said, uses 'the constitutive question' as I do, but others – including Paul Boghossian (1989: §7) – describe the relevance question as 'constitutive', in order to clarify that this latter question is not epistemological. My presentation of the constitutive question in this article does not imply that Boghossian, or any other philosopher who follows him in his use of 'constitutive', is addressing the constitutive question as I describe it.

this level that philosophers grappling with it have found themselves wading into choppy waters concerning the nature of objectivity and the mind-independence of the facts that, according to some, make up the world.[2] The constitutive question would not have such implications, if it were concerned merely with how our words come to express the concepts that they do.

To see what I mean, note that a satisfactory answer to the constitutive question must allow for our rules to 'issue their requirements independently and in advance of our appreciation of them' (Wright 2007: 482; see also Wright 1980: 21). John McDowell describes this feature of meaning as its 'normative reach' (1992: 273-274). The meaning – the concept – we attach to a word reaches ahead of our actual use of the word in the following sense: given the way the world has developed, the concept determines what is to count, in any circumstance, as its correct application in advance of anyone's actually being in that circumstance. Metaphorically speaking, once we understand a concept, its pattern of correct application is laid out ahead of us.

At the beginning of his magisterial early study of Wittgenstein's rule-following considerations (1982), McDowell points out that if, in answering the constitutive question, we leave no room for our concepts having normative reach, we thereby abandon our 'familiar intuitive notion of objectivity': the idea that things are thus-and-so 'anyway, whatever the outcome of any …

[2] For such a view of facts, see Wittgenstein 1922 and McDowell 1994.

investigation' (1984: p. 222). For such objectivity to be in place, there must be ways that things could be *correctly be said (or thought) to be* anyway, whatever we in fact go on to say (or think) about them; and this just requires that the concepts our words express have patterns of correct application that extend independently of how we actually go on to apply them (1984: 222). McDowell is right to think that repudiating meaning's normative reach and, with it, our familiar concept of objectivity, is to endorse 'a kind of idealism' (1984: 222). It is the familiar empirical idealism that has it that the facts of the matter in any given circumstance are determined by our verdicts or, more broadly, by our 'reactions to the case' (Wright 1989: 304).

That both McDowell and Wright see Wittgenstein's rule-following considerations as having the potential to nudge us towards idealism (McDowell 1984: 325; Wright 1982a: 252; 1982b: 246-247) demonstrates that they have the constitutive question in mind here. Neither the epistemological nor the relevance question could raise this particular metaphysical spectre. And yet, intriguingly, having explicitly cautioned against sliding towards idealism, McDowell says this towards the end of his paper:

> When we say '"Diamonds are hard" is true if and only if diamonds are hard', we are just as much involved on the right-hand side as the reflections on rule-following tell us we are. There is a standing temptation to miss this obvious truth, and to suppose that the right-hand side somehow presents us with a possible fact, pictured as an unconceptualized configuration of things in

> themselves. But we can find the connection between meaning and truth illuminating without succumbing to that temptation. (McDowell 1984: 255)

According to McDowell, then, Wittgenstein has shown us that the fact that diamonds are hard is not 'an unconceptualized configuration of things in themselves', but something in which we are 'involved'. But what could such 'involvement' be? Presumably, given the intended contrast with the notion of a fact that McDowell rejects, it consists in this fact about diamonds being in some way grounded in facts about us. And since this consequence is supposed to follow from Wittgenstein's reflections on rules and rule-following, McDowell's thought would seem to be something like this: what it is for diamonds to be hard is explained, at least in part, by how we carry on, that is to say, by contingent truths about our linguistic practices and ways of thinking. Saying this takes us beyond the mere truism that we use language to express the said fact and towards the idealist-sounding thesis that the fact that diamonds are hard is a 'precipitate of language' (Gaskin 2019: 1325).

No doubt, the threat here is an idealism distinct from the empirical kind involved in rejecting meaning's normative reach altogether: it is not being denied that the pattern of correct application we grasp, in grasping a concept, extends to the relevant case independently of the outcome of any particular investigation. Such a pattern, on McDowell's view, is ratification-independent, while not 'extend[ing] *of itself* to new cases' (1984: 256). And yet it seems to

be the case, according to McDowell's Wittgenstein, that facts about our reactions and attitudes partly explain both *what it is* for diamonds to be hard and *why it is correct*, when developing a numerical series by repetitively adding 2, to continue the series '1000, 1002, 1004, 1006 …', and not '1000, 1004, 1008, 1012 …'. Even if 'our reactions and attitudes' are not merely those of English speakers, but include those of any beings humans can recognize as language users, we will still want to insist that they have no bearing on the facts concerned. Any temptation to think that these facts are partially explained by the fact we are 'so minded' (to use Jonathan Lear's phrase (1982: 385-386)), ought to be resisted.[3]

2. 'Hard platonism'

In contrast to the 'quietism' explicitly defended by him in his later work (e.g. McDowell 1991; 1992; 1994; 2009), the McDowell of his 1984 paper is drawn into what presents itself as a substantive philosophical dispute over how a rule's standard of correct application is determined. This dispute assumes that

[3] Lear himself does not put up enough resistance. He says that '[a]fter studying the later Wittgenstein we are tempted to say that ['7 + 5 equals anything at all, just as long as everyone is so minded'] expresses some sort of truth'; but he goes on to say that is not an *empirical* truth because 'the notion of people being "other-minded" is not something on which we can get any grasp' and so does not 'for us' express a real possibility (1982: 386). Suggesting that the idealism involved here is something we cannot (fully) represent to ourselves is no comfort to those of us with a robust sense of reality.

the constitutive question is one that calls for some constructive, as opposed to therapeutic, philosophy. On this way of thinking, shared between the dispute's participants, how a rule can come to impose genuine requirements on us is mysterious, and we must deploy some philosophical theorizing in order to resolve the mystery. (Of course, McDowell does not say this himself, but, as we shall see presently, what he does say commits him to it.)

A perceived enemy here is what is commonly labeled 'platonism' about rules (Child 2011: 124-126), but which I shall call 'hard platonism'. Hard platonism has it that a concept's pattern of correct application is supplied by bare reality in itself, so that continuing our series '1000, 1002, 1004, 1006 …' is determined as correct because this way of doing it 'is absolutely the *simplest* or *most natural*' (Child 2011: 124). But while this conception clearly gestures at accounting for meaning's normative reach, McDowell alleges that it does so at the considerable cost of making rule-*following* uncanny and obscure. Learning how to extend a numerical series, for example, is disfigured into a matter of our somehow 'get[ting] our practice of judging and speaking in line with the rule's impersonal dictates' (1984: 231); and both he and Crispin Wright argue that such a picture fetishizes understanding as an item's coming before the mind that mysteriously – magically? – represents whether any possible application is correct or not (Wittgenstein 1953: §188; Wright 1989: 300-302).

I find this objection to hard platonism unpersuasive. But more on this later. What is striking is how tempting it is to suppose that avoiding hard platonism requires us to go into just the same line of business as the hard platonist, namely, answering the constitutive question by engaging in constructive philosophy. So we see McDowell arguing that we must make out a 'middle position' that repudiates hard platonism, but does so without slipping into the empirical idealism that we have seen abolish our concepts' normative reach altogether:

> What Wittgenstein's polemic … makes untenable is the thesis that possessing a concept is grasping a pattern of application that extends *of itself* to new cases… There *must* be a middle position. Understanding is grasping patterns that extend to new cases independently of our ratification, as is required for meaning to be other than an illusion (and – not incidentally – for the intuitive notion of objectivity to have a use); but the constraints imposed by our concepts do not have the platonistic autonomy with which they are credited in the picture of the super-rigid machinery.[4] (McDowell 1984: 256)

Evidently, in this passage McDowell is gesturing at what he takes to be a distinct region of philosophical space that lies between the abandonment of meaning's normative reach and the idea that a concept's pattern of correct application extends 'of itself' to new cases. If he is right, our meanings have a

[4] This tendency to think of the understanding as the operation of a 'super-rigid' machine, in which we engage with the 'rails' laid down by our meanings, is discussed in §3 below.

normative reach that is in some way conditioned by facts about how we find it natural to carry on.

It is hard to see how he *can* be right, though. McDowell's characterization of the flaw in hard platonism – namely, its taking our rules' patterns of correct application to extend *of themselves* to new cases – imposes a task on him to spell out exactly how the way in which such patterns extend is, at least in part, shaped by our practices. We must not suppose, McDowell says, 'that meanings [i.e. concepts] take care of themselves, needing, as it were, no help from us' (1984: 254). The 'as it were' here registers that the talk of 'help' is metaphorical, of course; but the metaphor's aptness lies in its figuratively signalling a commitment to the idea that our practices play a role in determining the progress of a concept's pattern of application to new cases. In short, *we*, *via* our practices, play a part in fixing the dictates of the concepts we grasp: facts about us figure in an explanation of why our concepts' patterns of correct application extend as they do.

An unsatisfying feature of McDowell's discussion is that he takes us this far, but no further: he points us in the direction of a place we can seemingly reach only by doing constructive philosophical work, but – no doubt, due to his official quietism on the matter – appears to opt out of doing this work himself. But much more seriously, there is a problem even conceiving of how the kind of middle position McDowell gestures at could ever satisfy us. If we

'help' our rules' patterns of correct application extend to new cases, then the facts that make up the world – for example, facts concerning the correct development of numerical series, as well as empirical facts such as diamonds being hard – lose their objectivity: they become dependent upon, grounded in, subjective contingencies: how we find it natural to carry on, what we see as doing the same thing, and so on. And we rejected this idea at the end of §1.

All of this suggests to me that McDowell gets off on the wrong foot – at least in his early work (1984) – by thinking, albeit implicitly, of the constitutive question as demanding a substantial answer.[5] Wittgenstein's reflections on

[5] If we read McDowell's subsequent, more thoroughgoing, quietism into his 1984, we might be tempted to regard him as engaging merely with the relevance question, and not with the constitutive question in the way in which I here describe. On such a reading, his claim that we are 'involved on the right-hand side' of a biconditional such as '"Diamonds are hard" is true if and only if diamonds are hard' means, not that we are involved in the explanation of what it is for diamonds to be hard, but merely that our linguistic practices determine that the sentence 'Diamonds are hard' expresses the proposition that diamonds are hard.

But such a reading makes no sense of McDowell's dialectic in the paper. As I explained in §1, his claim that *idealism* would be a threat if the pattern idea were dropped, and his explicit determination to seek 'a middle position' between such idealism and hard platonism, would be inexplicable were he not addressing the constitutive question as I have formulated it.

Furthermore, the thought that, for McDowell, our involvement on the right-hand side of the aforementioned biconditional means anything less than that we are involved in what it is for diamonds to be hard is falsified by the matter that both immediately precedes and follows this remark. Four paragraphs before he introduces the biconditional, McDowell says that Wittgenstein's rule-following considerations 'attack a certain familiar picture of facts and truth': specifically, the idea that 'a genuine fact must be a matter of the way things are in themselves, utterly independent of us' (1984: 254). And directly after he says that we are involved on the biconditional's right-hand side, he makes the observation that we should not think of the biconditional's right-hand side as presenting us with 'an unconceptualized

rule-following, properly understood, do, indeed, tell us that what we mean needs *some kind* of 'help' from us; they do, indeed, demonstrate that we are 'involved' in *some way* on the right-hand side of a biconditional such as '"Diamonds are hard is true" if and only if diamonds are hard'. But the nature of the 'help' we provide to the things we mean – the nature of our 'involvement' in the facts – raises no threat of idealism. This is because answering the constitutive question correctly does not take us into the rich, metaphysical territory gestured at by McDowell's talk of finding a 'middle position' between empirical idealism and hard platonism.

This conclusion will, I expect, be broadly welcomed by McDowell's later, more resolutely quietist, self (e.g. McDowell 1992, 1994, 2009). But, as we shall see in §5, there are two significant respects in which my 'diagnostic deconstruction' of the felt need for constructive philosophy in this area (McDowell 1992: 278) departs from his. First, I argue that if we are to cut the constitutive question down to size, we must do something that McDowell fails to do: focus squarely on the pervasive, yet altogether misconceived, idea that

---

configuration of things in themselves' (1984: 254). The clear implication here is that we should seek to do justice to the pattern idea in a way that repudiates the hard platonist's conception of possible facts. This proposal would come like a bolt from the blue, were McDowell's conception of our involvement on the biconditional's right-hand side interpreted along the alternative lines suggested. It is intelligible only if he is offering an account of the fact expressed by the biconditional's right-hand side that is intended to replace the hard platonist's, one which takes this fact to be in some sense determined by how things are with us.

something external to the rule itself must be involved in determining the standard for its correct application. Only once we loosen the grip that this latter doctrine has on us does the particular sea of philosophy that I am primarily concerned with – in which we are seemingly forced to choose between empirical idealism, hard platonism, or a foggy 'middle position' – eventually subside.

Second, but relatedly, loosening the grip that this defective picture of how a rule's requirements are determined is of a piece with one use of a metaphor that Wittgenstein himself introduces, but to which he has a somewhat complicated and equivocal relationship: that of the rules we follow as 'rails invisibly laid to infinity' (Wittgenstein 1953: §218). McDowell thoroughly deprecates this imagery (1984: 231-232). Following, I believe, Wittgenstein's own lead, I do not. Indeed, thinking of the rules we follow along these lines can actually help us to stay on an appropriately quietist path with respect to the constitutive question.

Most commentators associate the aforementioned 'rules-as-rails' imagery unambiguously with the hard platonism that McDowell rejects: the myth-eaten and fetishized conception of meanings as denizens of 'a platonic realm', completely sealed off from how we find it natural to go on, once trained (Child 1991: 168). But in my view, thinking of rules as rails can, in fact, help us establish the credentials of the kind of quietist answer to the constitutive

question that the later McDowell favours. Hence, the position I advocate, inasmuch as it merits the label 'platonism' by virtue of its visualization of rules as rails, can be called 'quietist platonism'. Those (over-) familiar with the literature might regard this name as an oxymoron, but it is not.

3. The subjectivity of rule-following

Let me briefly focus on the Wittgenstein's treatment of the epistemological and relevance questions. Doing so will provide some useful background when I go on to characterize the answer to the constitutive question I want to defend.

How, at any given point, do I know what the requirements of the rule for keeping adding 2 are? This question can seem troubling because, as someone pressing the relevance question points out, there are infinitely many different ways of continuing the series beyond a certain point that agree with each other up to that point (Wittgenstein 1953: §185). So imagine that I know every step in the series until 1000. How do I know that I should continue the series, '1000, 1002, 1004, 1006 …', and not '1000, 1004, 1008, 1012 …'? The way in which I have developed the series up to 1000 is consistent with either way of continuing it; there is a rule according to which continuing in the latter style would count as doing the same thing as before.

Famously, Wittgenstein rejects the response to this question given by what we might call 'the interpretive model of rule-following'. According to the

interpretive model, understanding an expression consists in having an item come before one's mind that supposedly serves as 'a rule determining the application of [the] rule' (Wittgenstein 1953: §84): something that Wittgenstein calls an 'interpretation' (1953: §201). This item is supposed to be a representation of the rule's requirements: something that provides me with a reason for using the word in question as I do and to which I can appeal to justify my use at any stage. Applied to the now familiar case of obeying the instruction to keep adding 2, the suggestion might be something like this: a mental image or representation of the form '{2n}' tells me how to develop the sequence in accordance with the order to keep adding 2: upon reaching 1000, I can consult this representation, which, in effect, tells me to carry on '1000, 1002, 1004, 1006 …'.

Wittgenstein is right that this will not do. Although understanding a word might sometimes *involve* an interpretation – for example, when one translates a word of a foreign language into one's home language in order to then figure out what the word means – this cannot be what such understanding *consists in*. This is because any such interpretation is on exactly the same footing as the sign whose meaning it supposedly represents. That is to say, if the question facing me is how to determine what the rule requires of me here, the very same question applies to the interpretation supposedly supplying this knowledge. For this latter thing is no more intrinsically meaningful than the

sign itself (McGinn 1984: 17): *any* interpretation, such as the representation '{2n}' in our example, admits of being taken in any number of ways. And so, if the interpretive model is correct, and if rules must be interpreted in order to be grasped, then the interpretation itself stands in need of interpretation, and so the interpretive model launches us upon an infinite, and vicious, regress. The moral here is that interpretations are categorially incapable of doing the job they have been dreamt up to do. An interpretation does not fix the application of a rule; interpreting a rule just substitutes one expression of it for another (Wittgenstein 1953: §201). For an interpretation to help us here, we must already be able to do the thing that it supposedly makes possible.

This is very familiar and, as philosophical interventions go, relatively uncontroversial. Its relevance for me lies in how fully appreciating it enables us to diagnose as illegitimate one particular use to which the 'rules-as-rails' imagery might be put. Gripped by the idea that grasping a rule can only consist in interpreting it, yet dimly aware of the devastating effect of the regress of interpretations that this idea triggers, someone might seek to hold on to the interpretive model while trying somehow to bring this regress to a halt. It is this strategy that Wittgenstein has in mind when he says that someone committed to the view of understanding as essentially interpretive will want to posit an interpretation that 'mustn't be capable of interpretation' but is 'the last interpretation' (1969: 34). The 'rules-as-rails' metaphor can be used in an

attempt to elaborate this idea in the following way: what makes it the case that the concept we grasp has the pattern of correct use that it has is that reality supplies this pattern as absolutely the most natural or simple (Child 2011: 126); and the interpretation of the rule that is itself uninterpretable is that which transparently and unequivocally presents this pattern.

This, it seems to me, is how Wittgenstein primarily envisages someone appealing to the metaphor of rules as rails laid to infinity. He asks

> [w]hence the idea that the beginning of a series is a visible section of rails invisibly laid to infinity? Well, we might imagine rails instead of a rule. And infinitely long rails correspond to the unlimited application of a rule. (1953: §218)

And his answer is that this way of thinking is part of an (inevitably doomed) attempt to come up with a last, uninterpretable interpretation of the kind just alluded to. Grasping such an interpretation, on this view, is an 'act of meaning [that] in its own way already traversed all those steps' constituting the relevant pattern of correct use: one's mind, so to speak, flew ahead and took all of the steps instantaneously (Wittgenstein 1953: §188). Since these steps were all taken in the act of grasping the interpretation, there exists for us no *choice* in how to apply the word (Wittgenstein 1953: §218). Understanding, according to this picture, is like the operation of a 'super-rigid' machine (Wittgenstein 1953: §197): the interpretation concerned provides a kind of guidance that

exactly bridges the gap between past and future use (McDowell 1984: 230). What it says, unlike other interpretations, cannot be variously interpreted.

Clearly, this is a fantasy. Even if we accept that it makes sense to say that there is an absolutely correct way to apply a rule, delivered by the standards of reality, there could not be an interpretation that could only be taken as representing such a standard. This is because *any* interpretation, however extravagantly described, can itself be variously interpreted; merely calling an interpretation 'the last interpretation' does not alter this fact. An interpretation, however we describe it, is in itself a lifeless piece of syntax (McGinn 1984: 17). So the appeal to the 'rules-as-rails' picture to try to prop up the interpretive model of rule-following is hopeless.

What this means is that the aforementioned 'blind' rule-following must be possible. We must be able to follow rules *neat*, without interpreting them. It is at this point that a familiar Wittgensteinian story heaves into view. What makes 'blind' rule-following possible, Wittgenstein explains, is our having been initiated into a custom (1953: §199): a nexus of practices in which we are trained to use words in specific ways and, in the course of so doing, come to see, quite unreflectively, certain ways of applying words, and not others, as natural and as 'going on in the same way' as before. Once this conceptual framework becomes engrained in us as a part of our 'second nature' (McDowell 1994: 95), the apparent need to mediate between a rule and its pattern of

correct application vanishes: we just apply the word as we have been trained to do (McDowell 1984: 239). On being told to keep adding 2, for example, we know that we should continue the series '1000, 1002', 1004, 1006 …'. We see this, and this alone, as doing the same thing as we did before we reached 1000; given our training, the question of whether some other way of developing the series is equally in line with what we did before simply never presents itself to us.

It remains true that any way of applying a word in future is compatible with its use up to now, *on some reading or other* (Wittgenstein 1953: §201). But there is only one pattern of use that leaps out at us as correct, given our immersion in the nexus of practices within which we have been trained in our use of words. Had our linguistic practices been different, we might have seen continuing the series '1000, 1004, 1008, 1012 …' as carrying on in the same way as before. In such a circumstance, we would have followed different rules: our words would have expressed different concepts, and so we would have meant something else by 'keep adding 2'. Call this dependence of which rules we follow upon contingencies in our sensibility – in essence, a sketch of an answer to the relevance question – 'the subjectivity of rule-following'.[6]

[6] I call this a 'sketch' of an answer to the relevance question advisedly. There is a good deal more more to say here. McDowell (1984; 1994: 92-05) does much to fill out this picture.

4. ‘Rules-as-rails’ imagery and the constitutive question

Metaphors, similes, and pictures have a place in philosophy, and particularly so for Wittgenstein. Sometimes they hold us captive (Wittgenstein 1953: §115), committing us to misleading analogies and preventing us from seeing other possibilities. But equally, thinking figuratively about something might remind us of something important or help us to appreciate its significance. For example, in describing the aim of his philosophical method as being that of ‘show[ing] the fly the way out of the fly-bottle’ (1953: §309), Wittgenstein elegantly combines two convictions central to his philosophical approach: that someone in the grip of one of philosophy’s perennial problems is not so much grappling with something deep that has to be *solved*, as struggling in a trap from which they have to be *released*; and that such traps, like fly-bottles in their transparency, are difficult to recognize as such from the inside. The metaphor works by presenting us with an image that jolts us into reconsidering our philosophical predicament, motivating a philosophical method that, on a case-by-case basis, sensitively seeks to break down the conviction that philosophical theory building is necessary.

A single metaphor, simile, or picture might be helpful or profound if used in one way, less so, or even misleading, if used in another. This is certainly true of talk of rules as ‘rails’. We saw in §3 that it is fruitless to use this metaphor as part of an attempt to prop up the interpretive model of rule-following: no

interpretation can be 'final' in the sense of admitting of being taken in only one way. In the wake of this fact, introducing the idea that rules are rails laid to infinity, and coupling it with a conception of understanding as a 'super-rigid' engagement such rails, is just obfuscation. This, I think, is largely what Wittgenstein is getting at when, after having presented us with such a model of rule-following, he asks this: 'But if something of this sort really were the case, how would it help me?' (1953: §219). This question is rhetorical, of course: the picture in play here is a 'philosophical superlative' (Wittgenstein 1953: §192), mere sleight of hand in the face of a ground-floor truth about the multiple interpretability of interpretations. Such an item cannot deliver on the promise of showing someone what a rule demands of her.

Intriguingly, however, Wittgenstein does not dismiss *tout court* the idea that it can be insightful to picture rules as rails. For after he has rejected an appeal to this metaphor to prop up the interpretive model of rule-following, he goes on to acknowledge that thinking of rules as rails nonetheless makes sense if 'understood symbolically'; and he elaborates what he means by this by saying, 'I should say: *This is how it strikes me*' (1953: §219). What Wittgenstein has noticed here is that conceiving of one's grasp of a rule in terms of being engaged on set of rails that, so to speak, binds one to a pattern of future use does justice to the phenomenology of 'blind' rule-following (1953: §219). It nicely captures the feeling of unquestioning security one feels when one follows a rule with

complete assurance, without feeling any need for an inner, mental justification of the sort presumed necessary by the interpretive model.

I believe that 'rules-as-rails' talk can be analogously enlightening when it comes to the constitutive question, too. A fully satisfying position on rules and rule-following will couple a recognition of rule-following's subjectivity with a commitment to the rules we follow being sufficiently objective to allow these rules to have normative reach. Thinking of rules as rails can help us to pull this off.

But only if the metaphor is applied in the right way. To be sure, we cannot achieve the desired equilibrium between the subjectivity of rule-following and the objectivity of the rules we follow, if we deploy the picture of rails laid to infinity as an encapsulation the kind of the hard platonism elaborated in §2. Hard platonism attempts to provide what in this context counts as a *substantial* answer to the constitutive question: one that presupposes that there exists an explanatory 'gap' between a rule and its requirements, and which seeks to bridge this gap by explaining what, beyond the rule itself, determines these requirements (Wright 1989: 304). The answer that the picture of rules as rails supplies, once put into such service, is essentially this: since following a rule for the application of a concept across different objects is a matter of engaging in a practice of *doing the same thing* in one's use of a word or expression (McDowell 1981a: 203), there exists a totally impersonal

standard for this, delivered by the nature of reality (Child 2011: 124). And this is flat-out false, as the subjectivity of rule-following demonstrates. By *our standards* – the standards of people with 'the perceptions of salience, routes of interest, feelings of naturalness in following a rule, etc. that constitute being part of a certain form of life' (Lear 1982: 385) – we carry on in the same way as before, in obeying the order to keep adding 2 beyond 1000, by continuing '1000, 1002, 1004, 1006 …'. But continuing the said sequence '1000, 1004, 1008, 1012 …' would count as doing the same thing as before, according to the standard of the unorthodox rule-follower Wittgenstein imagines (1953: §§185-187): someone occupying a form of life somewhat different to ours. Such a person would find it natural to understand this order as *we* would understand the order 'keep adding 2 up to 1000, but from then on, keep adding 4' (1953: §185). There is no fact of the matter as to which standard for *doing the same thing as before* is absolutely correct or the most natural. To the unorthodox rule-follower, her standard for doing the same thing is the natural, completely obvious one.

If we suppose that answering the constitutive question requires us to bridge an explanatory gap between a rule and the standard for its correct application, this question starts to look unanswerable. Hard platonism is untenable: unvarnished reality, we have just seen, cannot span such a gap. But equally, its polar opposite, empirical idealism, in its insistence that the gap

between a rule and its requirements in any given case is bridged by our reaction to that case, gives up on the idea of our meanings' having normative reach and, with it, our intuitive conception of objectivity.[7] This might make it seem that we have an obligation to explore the kind of compromise position that, his quietist inclinations notwithstanding, McDowell is drawn to in his 1984. This will be an account according to which the direction taken by a rule's correct pattern of application is determined by contingent facts about how *we* find it natural to go on, and what counts, for *us*, as doing the same thing, but in such a way that the rule has genuine normative reach. But as we saw in §2, the prospects for a middle position of this kind are dim. A consequence of such a stance is that worldly facts – concerning, for example, the hardness of diamonds or the correct way to develop a numerical series in accordance with the instruction to keep adding 2 – are in part determined by our (contingent) form of life; and we want to protest that such facts have nothing to do with this. They are facts that are to no extent down to us.

5. Rules-as-rails and 'quietist platonism'

[7] '[I]f the requirements of the rule are not constituted, as the Platonist thought, independently of our reaction to the case, what is there available to constitute them *but* our reaction? But that idea effectively surrenders the notion of a requirement altogether' (Wright 1989: 304)

McDowell's talk of an available 'middle position' between hard platonism and empirical idealism amounts to a double commitment. First, it goes along with the prevailing view that the constitutive question stands in need of a substantial answer in the sense introduced towards the end of §4: one which specifies whatever it is *external to the rule* that supposedly determines the rule's requirements. And, second, such talk proposes, albeit inchoately, a substantial answer that supposedly avoids the pitfalls of its two extremist competitors: one which takes a rule's standard of correct application to be determined by what we find natural, and what we count as doing the same thing, but in such a way that the rule's normative reach is preserved.

The later McDowell would, I think, be suspicious of his earlier self's adoption of a response to the constructive question that seeks to occupy the same explanatory space as hard platonism. Tellingly, he has come to regard his 1984 as being 'too hospitable' to the kind of reading of Wittgenstein that has the concept of *custom* figure in what he calls 'a constructive philosophical response to questions like "How is meaning possible?"' (1992: 275, n. 6). So I think he would now agree that the supposed gap between a rule and its pattern of correct application – that which motivates the quest for a substantive answer to the constitutive question – is an illusion. That is, I would expect him to concur with the suggestion that the constitutive question is one of those philosophical puzzles that has 'the *character* of depth' (Wittgenstein 1953: §111,

my italics), but which does not in fact call for any philosophical heavy lifting. And yet nothing in McDowell's writing since 1984 offers what we need at this point: 'a reminder for a particular purpose' (Wittgenstein 1953, §127): something that will show us exactly where someone in the grip of the kind of explanatory presumptuousness shared by the hard platonist, the empirical idealist, and his former self has gone wrong.

Some might find this judgement harsh. After all, does not McDowell explicitly reject the idea of a 'gap' between a rule and its requirements in remarks such as the following?

> [W]hat one responds to when one follows a signpost is not something that stands at a remove from a sorting of behavior into what accords with the rule expressed by the signpost and what does not, separated from the sorting of behavior by that supposed gap, a gap that "for a moment" (§201), when we embark on the train of thought that generates that puzzlement, we think could be bridged by an interpretation. (McDowell 2009: 368; see also 1992: 270-273)

The answer to this question is 'No'. McDowell's target here is 'the master thesis': the idea that a thing before someone's mind must be interpreted, if it is to sort extra-mental items into those that accord with it and those that do not (McDowell 1992: 270). And his wholly correct response is that once this thesis is rejected, as it should be, there is nothing for an interpretation to *do*, since, for the person properly embedded in the appropriate system of customs, the rule perfectly well sorts extra-mental items into those that do, and do not,

accord with it. But once things are put like this, it becomes evident that what McDowell says here does not so much as address the constitutive question: the question of what determines the *standard* for a rule's correct application. To repudiate the apparent 'gap' between a rule before the mind and one's grasp of its requirements – that which is opened by 'the master thesis' and which we are tempted to think could be bridged by an interpretation – is not to repudiate the wholly different explanatory 'gap' that supposedly exists between a rule and the determination of these requirements. The former is epistemological, concerning how we can know what a rule requires of us, while the latter is constitutive, concerning how the rule's requirements are themselves determined. Indeed, it would be a striking *non sequitur* to think that this latter 'gap' – that which invites the adoption of hard platonism or empirical idealism – could be closed by rejecting 'the master thesis'. How could pointing out that rules can be grasped non-interpretively help us decide how the requirements of these rules are fixed?

True enough, McDowell has come to present Wittgenstein's views on meaning and understanding as an exemplar of the kind of 'naturalized platonism' about normativity that he favours (1994: 92); and, it is part-and-parcel of this reading that '[t]here is no need for constructive philosophy, directed at the very idea of norms of reason, or the structure within which meaning comes into view, from the standpoint of the naturalism that threatens

to disenchant nature' (1994: 95). But my point remains. In order to make the constitutive question manageable, we need a specific, responsive diagnosis, of a kind unconsidered by McDowell, that makes plain what someone gripped by the urge towards constructive philosophy in this precise context has failed to see. And as I shall now explain, adopting the 'rules-as-rails' imagery – imagery that McDowell abhors (1984: 231-232) – can help in this process, enabling us to appreciate what is wrong with the way of thinking that makes a substantial answer to this question seem obligatory.

To recall our thinking from running in the *particular* groove that makes it look as if we need a substantial answer to the constitutive question, we must focus on the relation between a rule and its requirements. So what determines the standard for a rule's correct application? Not some other thing or things, such as our linguistic practices, or reality 'in itself', as the hard platonist thinks of it, but just this: the rule. In assigning a meaning to a word or expression, we establish a rule for using it correctly, and what makes this the rule that it is is its particular pattern of correct use. The rule *just is* a certain such pattern. This is the answer to the constitutive question, and it is this that someone who thinks that this question calls for some constructive philosophy needs to be reminded of.

The mistake, then, is to think that a rule's requirements have to be established by *something else*: that there must be something external to the rule

that mediates between itself and its pattern of correct application. Once we fall into this error, we throw ourselves into the apparently substantive, yet ultimately stultifying, project of trying to figure out what this mediating factor could be. Since neither reality in itself, nor our reaction to the case at hand, can be said to fulfil such a mediating role, we are led inexorably to the conclusion that there must exist a 'middle position' between these two extremes. But as we have seen, such a position, insofar as it can be made out at all, cannot give us what we want: in treating worldly facts as somehow grounded in facts about our form of life, it violates our sense that how the world really is has nothing to do with contingencies concerning what we find natural, or what we count as carrying on in the same way.

The way out lies in appreciating that the task seemingly demanded of us here is spurious. A rule's normative reach is *internal* to it in the sense that its associated standard of correctness does not stand in need of being founded on anything else. The rule's requirements are, truistically, those of the rule itself: the meaning we have attached to the word or expression. It seems to me that a further way in which the picture of rules laid to infinity can help us is by enabling us to internalize this minimalist story. First, visualizing a rule as a set of rails laid to infinity serves as a reminder that the rule has genuine normative reach: the pattern of application we grasp when we understand a concept does, indeed, extend, independently of the outcome of any investigation, to the

relevant case (McDowell 1984: 222). And this imagery, if appreciated in the right spirit, also makes us appreciate the key insight I have been driving at: namely, that there exists no explanatory gap between a rule and its requirements that, so it seems, must be filled by something external to the rule itself. For according to the 'rules-as-rails' picture, a rule *is* a set of rails laid to infinity: there is *nothing more* to the rule than its pattern of correct use. So the gap that the empirical idealist and the hard platonist see themselves as trying to bridge is illusory. The metaphor of rules *as* rails thus nicely points out, in an engagingly figurative way, that the kind of explanation these rival positions seek to give is uncalled for: unavailable, in fact.[8]

Since the 'rules-as-rails' picture is standardly described as 'platonist', the position I have set out, by dint of using this imagery, is aptly thought of as

[8] Wright, for one, is unhappy with such quietism on the constitutive question. According to him, if we think of a rule in the way recommended here – as something that *in itself* lays down a ratification-independent pattern of correct application, so that 'there cannot be a legitimate puzzle *how* a rule does that' (2007: 483) – then the problems that we seemingly avoid just pop up elsewhere. Specifically, his charge is that adopting such quietism merely invites the following, dizzyingly intractable question: 'what makes it possible for there to *be* such things as rules, so conceived, at all?' (2007: 483). How, in other words, can a rule *we* grasp 'mandate (or allow or forbid) determinate courses of action in an indefinite range of cases that its practitioners will never have explicitly considered or prepared for?' (2007: 483)

But the anxiety that Wright gives voice to here is uncalled for. In an echo of something McDowell says about intention, it is innocuous, and certainly nothing that calls for substantial philosophical explanation, to say that rules 'contain within themselves the distinction between conformity and non-conformity' (1991: 315). Rules *just are* things that have their standards of correct application built in as standard. To think otherwise is to fetishize their normative reach.

platonism of a sort. But by contrast with the hard platonism discussed in §2, it is platonism of a distinctly quietist kind. There is no commitment either to a conception of the rules we follow – the concepts we express in words – as inhabitants of a platonic realm or to the idea that there is one way of continuing an arithmetical sequence that is absolutely the simplest or the most natural. (That the meaning *we* assign to an expression thereby sets up a pattern of correct use is quite compatible with the fact that the rules we follow are those that chime with *our* contingently conditioned sense of what counts as going on in the same way.) The normative reach of our concepts is real – a rule's pattern of correct application is determined in advance of our investigations – but there is nothing philosophically suspect about this. We do not have to engage in constructive philosophy to show how this is possible. A concept's standard of correct application is internal to it, and thereby assigned to the word that expresses that concept. That is it.

Whether the position I have set out is Wittgenstein's is moot and, perhaps, undecidable. Quietist platonism, in its insistence that some apparently substantive philosophical work is, in fact, unnecessary, coheres very nicely with Wittgenstein's officially therapeutic meta-philosophical position. A philosophical theory designed to bridge the gap between a rule and its requirements might, for him, be just one more 'house of cards' (1953: §118); and, accordingly, the proper philosophical task will be to expose it as such and

to free us from the temptation to try to dig below the level at which explanation comes to an end (Wittgenstein 1978: VI, §31). And yet Wittgenstein himself evidently on occasion suffers from this very temptation himself, notably when appearing to characterize his position in the following, strikingly Kantian terms:

> We have a colour system as we have a number system. – Do the systems reside in *our* nature or in the nature of things? How are we to put it? – *Not* in the nature of numbers or colours. (1981: §357)

This is by no means an endorsement of a crude empirically idealist answer to the constitutive question: Wittgenstein is clear that human agreement does not determine what is true or false (1953: §241). But in inching towards the conclusion that our number system and colour system reside in our nature, rather than in the respective natures of colours or numbers, he would seem to be anticipating McDowell's early 'middle position', gravitating towards the thought that *we* figure in explanations of what it is for the possible facts reported on the right-hand sides of biconditionals such as '"Diamonds are hard" is true if and only if diamonds are hard' and '"7 + 5 = 12" is true if and only if 7 + 5 = 12' to obtain. A.W. Moore describes this urge as an urge towards 'transcendental idealism' (Moore 1997: 135), a position which he convincingly

argues to be incoherent (1997: ch. 6). [9] But whether we think that such idealism is, as he believes, self-stultifying, or merely false, the point is that the constitutive question does not lead us into such a metaphysical minefield. Its correct answer is an acknowledgement that the rules we follow *are* their patterns of correct application, and hence that the relation between a rule and its requirements does not admit of the weighty, metaphysical explanation that seems to make a choice between empirical idealism, hard platonism, and transcendental idealism obligatory. Wittgenstein, I think, glimpsed this, but seems to have struggled to internalize it fully.

6. Misdirected criticism of the 'rules-as-rails' imagery

Quietist platonism says this: if we detach the 'rules-as-rails' picture from a commitment to hard platonism, using it instead to convey the thought that a rule's normative reach, since internal to the rule itself, is unproblematic, then this imagery is beneficial. It can serve as a reminder for a particular purpose (Wittgenstein 1953: §127): namely, that of uncovering something obvious

[9] McDowell has certainly felt the pull of such transcendental idealism. In his 1984 he says that the thesis that the community has no standard to meet 'may have some merit', provided it is conceived as 'an attempt to say something that cannot be said but only shown' (1984: 256). And in an earlier paper, he describes Wittgenstein, with evident approval, as having replaced realism with anti-realism at the 'transcendental level' (1981b: 342). (In a note added in 1997, he rows back from this latter description in a way that coheres with his later, more thoroughgoing, quietism (1981b: 342, n. 68).)

which, when recollected, can dissolve our misconceived sense of obligation to find something external to a rule to determine its requirements.

Some philosophers, however, think that such imagery can only ever be mystery-mongering and unhelpful. Wright, for one, believes that thinking of rules as rails laid to infinity cannot but commit its user to an account of rule-following dripping with 'cognitive pretentiousness' (1989: 301; see also McDowell 1984: 230-232). According to him, grasping rules-as-rails can only be pictured as 'a kind of hyper-cognitive felicity' by which one magically leaps to 'an informational state by which the accord or clash with the rule of any proposed move is settled' (Wright 1989: 300).

This allegation does not stick, though. The mythical 'informational state' here is the supposedly uninterpretable *interpretation* rejected as a fairy tale in §3: a mental item that is supposed, incredibly, to contain within it a word's whole use and which, when it comes before the mind, makes meaning something by a word a matter of flying ahead and taking all the steps in its pattern of correct application instantaneously (Wittgenstein 1953: §188). We may agree with Wright that there can be no such thing, since *any* interpretation is in itself semantically inert (Wittgenstein 1953: §§ 191-192). It is the failure to see this that is the source of the interpretive model of understanding's incurable travails.

But why think that the quietist platonist's use of the 'rules-as-rails' metaphor commits her to such an account of understanding? If we distinguish (as we should) the question of what it is to follow a rule from that of the kind of objectivity these rules enjoy – then the following possibility opens up. The quietist platonist can agree with everything Wittgenstein says about rule-*following*: specifically, she can say that when one grasps a rule, its ratification-independent pattern of correct use is present to one only in the sense in which one has gained the *ability* to trace this pattern, and not in the 'strange' sense according to which one has an uninterpretable interpretation before one's mind which somehow (and miraculously) determines one's future use. One's future use is, *of course*, present in *some* way when one understands a word (Wittgenstein 1953: §195), but the way in which it is present – not via a miraculous 'last interpretation', but as an ability to use the word in accordance with one's training, utilizing a shared conception of *doing the same thing as before* – is compatible with thinking of the rule itself as a set of rails, as the quietist platonist recommends.

7. Our 'involvement' in the facts, revisited

The relation between a rule and its requirements is determined neither by our contingent form of life, nor, as the hard platonist supposes, by the existence of

an absolute and impersonal standard for carrying on in the same way. Each of these answers to the constitutive questions tries to dig too deep.

The thing we need to be reminded of, and which thinking of a rule *as* a rail laid towards infinity can help us recall, is that a rule *is* a certain pattern of correct use. A rule's normative reach is not something that has to be determined by something *external to the rule*; a rule's pattern of correct application does not stand in need of being secured by some other thing. True enough, which rules are the ones we follow is explained by the myriad things falling under the heading of our 'form of life' (Wittgenstein 1953: §241): what leaps out at us as carrying on in the same way will be determined by facts about our biological nature and the nexus of interests, activities, and practices that define us as the social beings we are. But this subjectivity figures in the subjectivity of rule-*following*; it does not compromise the objectivity of the rules we follow.

So where does this leave our 'involvement' on the right-hand side of a biconditional such as '"Diamonds are hard" is true if and only if diamonds are hard'? In a safe place. That we use the concept 'hard' to characterize the world *at all* is a fact contingent upon our form of life: had our nature and practices been different, we would not have possessed such a concept; we would not have thought in those terms (Moore 1997: 133). Accordingly, the fact that our meanings do not 'take care of themselves, needing, as it were, no help from us' (McDowell 1984: 254), admits of a similarly deflationary reading. The 'help'

we bring to what we mean is that of determining *which rules are expressed by our words*; and this we do by bringing to bear our subjectively conditioned criteria for carrying on in the same way, thereby making it possible for us to follow rules 'blindly'. What we do *not* do is play any sort of role in determining how our concepts' patterns of correct application extend to new cases. Our 'involvement' on the right-hand side of '"Diamonds are hard" is true if and only if diamonds are hard' thus amounts to no more than this: these concepts are ours because we happen to be the people we are. This leaves our concepts' normative reach and, with it, our familiar conception of objectivity secure and, indeed, unremarkable.[10]

University of Leeds, UK

References

Boghossian, P. (1989) 'The Rule-Following Considerations', *Mind*, 98: 507-549.

[10] A draft of this paper was given at a workshop on linguistic idealism at the University of Liverpool, which Richard Gaskin kindly organised in June, 2019. Many thanks to everyone there who gave me such helpful comments. Thanks, too, to Alex Miller, two anonymous referees for *The Philosophical Quarterly* and, especially, to my friend, Michael Morris, who first made me see the profundity in the rule-following considerations back in the late-1980s.

Child, W. (1991) ‘Review of *Meaning and Interpretation* (ed.) C. Travis’, *Mind*, 100: 162-171.

— (2011) *Wittgenstein*. London: Routledge.

Gaskin, R. (2019) ‘From the Unity of the Proposition to Linguistic Idealism’, *Synthese* 196: 1325-1342.

Kripke, S. (1982) *Wittgenstein on Rules and Private* Language. Oxford: Blackwell.

Lear, J. (1982) ‘Leaving the World Alone’, *Journal of Philosophy*, 79: 382-403.

McDowell, J. (1981a) ‘Non-Cognitivism and Rule-Following’, in S. Holtzman and C. Leich (eds.) *Wittgenstein: To Follow a Rule*. London: Routledge. Reprinted in McDowell 1998a, 198-218.

— (1981b) ‘Anti-Realism and the Epistemology of Understanding’, in H. Parret and J. Bouveresse (eds.) *Meaning and Understanding*. Berlin: De Gruyter. Reprinted in McDowell 1998b, 314-343.

— (1982) ‘Criteria, Defeasibility, and Knowledge’, *Proceedings of the British Academy*, 68. Reprinted in McDowell 1998b, 369-394.

— (1984) ‘Wittgenstein on Following a Rule’. *Synthese*, 58. Reprinted in McDowell 1998a, 221-262.

— (1991) ‘Intentionality and Interiority in Wittgenstein’, in K. Puhl (ed.) *Meaning Scepticism*. Berlin: De Gruyter. Reprinted in McDowell 1998a, 297-321.

— (1992) ‘Meaning and Intentionality in Wittgenstein’s Later Philosophy’, in P. French, T. Uehling, and H. Wettstein (eds.), *Midwest Studies in*

Philosophy 17: *The Wittgenstein Legacy*. Notre Dame: University of Notre Dame Press. *Reprinted* in McDowell 1998a, 263-278.

— (1994) *Mind and World*. Cambridge, MA: Harvard University Press.

— (1998a) *Mind, Value, and Reality*. Cambridge, MA: Harvard University Press.

— (1998b) *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.

— (2005) 'The True Modesty of the Identity Conception of Truth', *International Journal of Philosophical Studies*, 13: 83-88.

— (2009) 'Wittgenstein's "Quietism"', *Common Knowledge*, 15: 365-372.

Moore, A.W. (1997) *Points of View*. Oxford: OUP.

Wittgenstein, L. (1922) *Tractatus Logico-Philosophicus*, trans. C.K. Ogden. London: Routledge and Kegan Paul.

— (1953) *Philosophical Investigations*, trans. G.E.M. Anscombe, P.M.S. Hacker and J. Schulte. Oxford: Wiley Blackwell, 2009.

— (1969) *The Blue and Brown Books*. Oxford: Blackwell.

— (1978) *Remarks on the Foundations of Mathematics*, 3rd edn, trans. G.E.M. Anscombe. Oxford: Blackwell.

— (1981) *Zettel*, trans. G.E.M. Anscombe. Oxford: Blackwell.

Wright, C. (1980) *Wittgenstein on the Foundations of Mathematics*. London: Duckworth.

— (1982a) 'Strict Finitism', *Synthese*, 51: 203-282.

— (1982b) ‘Anti-Realist Semantics: The Role of Criteria’, in G. Vesey (ed.) *Idealism: Past and Present*, 225-248. Cambridge: CUP.

— (1989) ‘Critical Notice of Colin McGinn’s *Wittgenstein on Meaning*’, *Mind*, 98: 289-305.

— (2007) ‘Rule-Following Without Reasons: Wittgenstein’s Quietism and the Constitutive Question’, *Ratio*, 20: 481-502.