Having Meanings in View: Groundwork for a Substantive Account of Rule-following

Arnaud Petit

A thesis submitted in fulfillment of the requirements for the
DPhil in Philosophy

Brasenose College
University of Oxford

Table of Content

Acknowledgements **iv**

Abstract **vi**

List of Abbreviations **vii**

Introduction **2**

1. McDowell on Following a Rule **7**

1.1 Introduction 7

1.2 McDowell and the Regress of Interpretations 7

1.3 Quietism and Anti-Reductionism 13

1.4 Objectivity and Guidance 23

1.5 Making Philosophical Claims 30

1.6 Philosophy Without End 38

1.7 Philosophy Like Botany 49

2. Merleau-Ponty on Meaning and Motivation **52**

2.1 Introduction 52

2.2 The Project of the *Phenomenology* 52

2.3 The Content of Perception 56

2.4 Perception and Action 66

2.5 The Richness of Perceptual Meanings 71

2.5.1 Intentional Richness and Conceptual Mastery 74

2.5.2 Intentional Richness and Basic Behavioural Capabilities 78

2.6 Motivation, Guidance, and Objectivity 87

2.7 A Layered Conception of our Engagement with the World? 91

2.8 The Unfinishedness of Phenomenology 100

3. Wittgenstein on Intentionality and Internal Relations **103**

3.1 Introduction 103

3.2 Intentionality in the *Tractatus* 103

3.3 Intentionality in *Philosophical Investigations* 115

4. Locating Bedrock; Laying the Groundwork **127**

4.1 Introduction 127

4.2 What Features of Signs are we Responsive to? 128

4.2.1 Objectivity and Mere Conventionality 128

4.2.2 Creatures of Ours 134

4.2.3 A Partial Reduction 139

4.2.4 Kinds of Meanings 145

4.2.5 Purposes 147

4.2.6 Basic Patterns 154

4.3 What is the Relation Between a Sign and our Response to it? 159

4.3.1 Motivation 159

4.3.2 Guidance 165

4.3.3 Normativity 169

5. Following a Rule: The Contemporary Debate **171**

5.1 Introduction 171

5.2 Anti-normativism: A Brief Overview 173

5.3 The Intentional View: Rules and Regresses 177

5.4 Ginsborg's Primitive Normativity 184

5.5 Guidance, Objectivity, and Epiphenomenal Understanding 188

5.6 Intelligent Action 197

6. Conclusion **200**

Bibliography **204**

Acknowledgements

I would like to sincerely thank my primary supervisor, Professor Bill Child, for his support and guidance over the past three years. Our meetings have been deeply stimulating; challenging at times, but always helpful. If there is clarity in the thoughts herein expressed, it is in part thanks to him. I can safely say that he has made me a better, more rigorous, thinker and, for that, I am indebted to him. I would also like to sincerely thank Professor Stephen Mulhall, who kindly accepted to be my secondary supervisor, despite the already late stage of my dissertation and the disruption of the pandemic. I have benefitted greatly from our exchanges. They have led me to nuance and refine my reading of many authors and have fed my philosophical thinking on many issues – both within and beyond the boundaries of the present work.

I am grateful to Professor Mark Wrathall for his comments on drafts of chapter 2. As I am grateful to Professors Hannah Ginsborg, Anandi Hattiangadi, Robert Myers, Alex Miller and Claudine Verheggen, for their comments on early drafts of chapter 5, which I presented some years ago at the annual congress of the Canadian Philosophical Association. My thanks also go to Dr Yuuki Ohta and Professor Adrian Moore, for their feedback on the work I submitted for my transfer and confirmation of status.

I am especially thankful to Dr Natalie Barber for her unwavering support and encouragements over the past two years. Had it not been for her help with proofreading, the present dissertation would have remained a mix of American and British English.

Finally, I am grateful to the Social Sciences and Humanities Research Council of Canada (SSHRC) and the W.L. Mackenzie King Memorial Scholarships for their funding.

Abstract

The present work is an attempt to lay the groundwork for an account of rule-following that is both philosophically substantive and partially reductive. It is philosophically substantive in contrast with the kind of discourses that exemplify a quietist approach to philosophy. It is partially reductive inasmuch as it proposes to reconstruct the phenomena of rule-following, linguistically or conceptually shaped meaning, and understanding out of materials that do not presuppose either, though without depriving itself altogether of properly intentional resources.

Ultimately, the present work aims to answer two questions: (i) which feature of signs are we *responsive to* in cases of rule-following? and (ii) what is the relation between these signs and our responses to them in such cases? It does so, not in sterile isolation from any historical anchorage, but through a critical engagement with the writings of three authors: McDowell, Merleau-Ponty and Wittgenstein. It is, we could say, an attempt at constructive philosophy which takes form in dialogue.

List of Abbreviations

BB	Wittgenstein, L. (1969). *The Blue and Brown Books* (2nd Edition). Oxford: Blackwell.
PI	Wittgenstein, L. (2009). [1953] *Philosophical Investigations* (Fourth Edition). Oxford: Wiley-Blackwell.
PR	Wittgenstein, L. (1975). *Philosophical Remarks*. Chicago: The University of Chicago Press.
RFM	Wittgenstein, L. (1983) [1956]. *Remarks on the Foundations of Mathematics*. Cambridge, MA and London, England: The MIT Press.
LWPP-I	Wittgenstein, L. (1990). *Last Writings on the Philosophy of Psychology, Volume 1*. Oxford: Blackwell.
TLP	Wittgenstein, L. (2001) [1961]. *Tractatus Logico-Philosophicus*. D.F. Pears and B.F. McGuiness (trans.). Abingdon and New York: Routledge.
Z	Wittgenstein, L. (1967). *Zettel*. Berkeley, Los Angeles, and London: University of California Press.

Having Meanings in View: Groundwork for a Substantive Account of Rule-following

Introduction

Ever since the publication of Kripke's *Wittgenstein on Rules and Private Language*, the phenomenon of rule-following has been at the center of intense debates within so-called analytic philosophy. Some of these debates have been in part exegetical and have sprung from questions such as: "what does Wittgenstein really aim to achieve in his discussion of rule-following?", "Is Kripke's reading of Wittgenstein exegetically faithful?", etc. But most – if not all – of these debates have also involved some reflections as to what it *is* to follow a rule. It is this constitutive question – and the philosophical perplexities and apparent paradoxes it makes salient – which mostly interest us here.

Now almost forty years later, debates persist as to how best to account for the phenomenon of rule-following. It is against that backdrop of continuing philosophical disputes that the present project has emerged.

Our aim, very roughly, is to lay the groundwork for a substantive account of rule-following – one that avoids the philosophical perplexities and apparent paradoxes that have so often plagued such attempts. More precisely, we aim to answer two questions, in some details: (i) which feature of signs are we *responsive to* in cases of rule-following? and (ii) what is the relation between these signs and our responses to them in cases of rule-following?

Obviously in answering these questions, we must not fall prey to the regress of interpretations – nor to any of its variations. But in our view, a satisfying answer to these

questions must also meet two further conditions, which we have dubbed the objectivity and guidance conditions.

Whether or not an action conforms to a rule – or *accords* with it – depends on the requirements of the rule. The objectivity condition stipulates that these requirements cannot collapse with what we take them to be. What, as a matter of fact, accords with a rule cannot simply be what we think accords with it. Otherwise, we could not talk of anything being in accord with the rule – or going against it. If we frame this in terms of meaning rather than in terms rules and requirements, we get something like this: whether or not something accords with a sign on the ground of its meaning has to be independent of our alleged awareness or appreciation of its meaning.

As for the guidance condition, it stipulates that, in cases of rule-following, the rule – or the sign – must somehow *tell us* how to act. Our awareness of the requirements of the rule – or of the sign's meaning – must play a role in *guiding* our action. The guidance condition thus implies that the requirements of the rule – or sign's meaning – must be in principle available to us, if following a rule is to be possible at all. They must be within reach of our cognitive capacities. They might not be up to us, but it must be possible for us to be receptive and ultimately responsive to them.

Now this is a purposely rough sketch: both conditions are meant to capture fundamental insights many share about rule-following. Our task, in the following chapters, will be to show how these two conditions can be met. But, as our investigations progress, it will also involve further refining the insights they are meant to capture. Our approach is far from dogmatic.

Still, we can already make a few observations. The objectivity condition will mostly impose a constraint on our answer to the first question. Allegedly, the feature of signs we are responsive to, in cases of rule-following, are the very features that determine or ground what accords with the signs. And so, if the objectivity condition is to be met, these features have to be features *of* the signs themselves. They cannot be features *we take the signs to be having* or *project onto the signs*; they must be somehow intrinsic to the signs. As for the guidance condition, it will mostly impose a constraint on our answer to the second question: ultimately, not any relation between a sign and our responses to it will do. It will have to accommodate the insight that the sign can *tell us* what accords with it, that we are appropriately equipped to grasp or understand what it is telling us, and that our response to the sign is guided by that understanding.

Furthermore, it is already clear that the objectivity and guidance conditions can easily come into tension with each other, the satisfaction of one threatening the satisfaction of the other. The tension boils down to this: the objectivity condition requires the requirements of a rule to not be constitutively dependent on what we take them to be. But the most obvious ways to flesh out that independence makes it utterly mysterious how the requirements could then be epistemically available to us – and ultimately guide us.

Our investigations unfold as follows. Our initial aim, in chapter 1 is to get clearer on what is required to escape the regress of interpretations. We argue, with McDowell, that the only way to escape the regress is to not embark on the regress at all – and that this is only possible if we acknowledge that signs, in themselves, need not be "normatively inert". We then attempt to show that this insight is ultimately distinct from – and independent of – two other strands of McDowell's view on rule-following: namely his anti-reductionism

and his quietism. We end up suggesting that it may be possible to fruitfully situate McDowell's insight about signs within the context of an account of rule-following that is both philosophically substantive and partially reductive.

To be partially reductive our account needs to reconstruct meaning and understanding that are linguistically – or conceptually – shaped, out of materials that do not presuppose such meaning and understanding. This involves, we argue, digging below the level of meaning and understanding McDowell takes to be bedrock, though not so low as to entirely deprive ourselves of intentional resources. In chapter 2, we suggest that a number of Merleau-Ponty's insights on perception and action can help us in that endeavour. We focus especially on the conception of perceptual content and on the notion of motivation he lays out in the *Phenomenology of Perception*.

Merleau-Ponty's insights also suggest something interesting: to the extent that our responses to signs, objects and situations are grounded on a bedrock of intentionality, they are best conceived as *internally* r*elated* to these signs, objects and situations. In chapter 3, we attempt to clarify that suggestion by relating it to Wittgenstein's evolving view on intentionality and internal relations.

In chapter 4, and in light of our critical engagement with McDowell, Merleau-Ponty and Wittgenstein, we finally attempt to answer our two questions: (i) which feature of signs are we *responsive to* in cases of rule-following? and (ii) what is the relation between these signs and our responses to them in cases of rule-following? We spend some time showing how our proposal meet both the objectivity and guidance conditions.

Finally, in chapter 5, we attempt to situate our proposal within the recent debate on the normativity of meaning. We discuss the failure of the so-called intentional view on

rule-following at escaping the regress of interpretations – and variations of it. We also compare and contrast our proposal with the account of rule-following recently put forward by Hannah Ginsborg.

Chapter 1

McDowell on Following a Rule

1.1 Introduction

In what follows, I attempt to show that we can preserve McDowell's major insight about the nature of signs and of rule-following – while rejecting the quietism it is often conflated with. To begin with, I argue that McDowell's answer to the regress of interpretations consist in abandoning the assumption that signs, in themselves, just 'stand there' and in recognising, on the contrary that they are the kind of things that are intrinsically meaningful and so impose normative constraints on us. Then, I argue that we must distinguish this insight from McDowell's quietism. I also argue that McDowell's quietism and his anti-reductionism about meaning and understanding are two distinct commitments. With that in mind, I briefly discuss how McDowell situates – and how he further fleshes out – his response to the regress in the context of a broader discussion of objectivity and guidance. Finally, I come back to quietism and how best to contrast it with constructive philosophy.

1.2 McDowell and the Regress of Interpretations

According to McDowell, the point of Wittgenstein's rule-following considerations is not to answer the regress of interpretations – that apparently threatens the very ideas of meaning and understanding –, but to loosen the grip of a problematic assumption about

meaning and understanding that gets the regress going in the first place. Ultimately, the regress is not a genuine and pressing problem; it does not call for a philosophical solution. There is no such solution anyway. At first sight, the regress appears unavoidable: after all, is it not the case that a sign, no matter the medium, can always be variously interpreted? A signpost (e.g. "a stop sign") might mean something to us, but what would prevent it from meaning something else entirely to others? And does this not show that our understanding the signpost, our grasping its meaning has to be somehow mediated by an interpretation? Considered in itself, the signpost does not mean anything, or so the thought goes. It does not impose any normative constraints on us (i.e. it does not specify how we should respond to it, what would accord with it and what would go against it). It is, as McDowell puts it "normatively inert" (McDowell, 2013, 81). Only the signpost *under a particular interpretation* could impose such constraints on us. But then, the worry immediately reappears: is it not conceivable that the interpretation might be variously interpreted too? At each step, one could provide a further interpretation, but that would only push back the problem, not solve it.

To be sure, the issue is not merely that we are never able to communicate to others what we mean by our words – and more generally how we understand overt signs. Of course, the regress *does* raise such an issue, as the example of the stop sign makes it clear. Once we realise that the sign can be variously interpreted, we might supplement it with written instructions in an attempt to specify how we interpret it, what meaning it has for us (e.g. "stop at the intersection"). But then, these written instructions too can be variously interpreted. That is to say, nothing prevents this particular string of letters – or to be even less committal, this particular arrangement of physical marks – to mean something else to

others. We could then be tempted to accompany the written instructions with explanatory gestures for example, but these gestures too would be ultimately open to interpretation. Nothing, it seems, can ever guarantee us that others will understand these signs as we do and so grasp the meaning we are trying to convey. Even repeated agreement in our responses to the stop sign could not give us the assurance of a shared understanding. Interpretations that accommodate these responses and yet ultimately come apart are always conceivable. And since, each interpretation just 'stands there', calling for a further interpretation, the problem remains. Therefore, if accepted, the regress does cast doubt on communication, on our ability to ever engage in interactions that are dependent upon the acknowledgment of common meanings.[1]

But the regress leads to an even more radical conclusion. If what a sign means is always dependent upon an interpretation and that interpretation is itself a sign which can, by its very nature, always be variously interpreted, then the regress ultimately threatens the idea that signs, in general, can have determinate meanings and be understood in particular ways. And this would imply that signs can never impose specific normative constraints on us. The apparent impossibility of communicating to others what overt signs mean to us is simply a symptom of that more fundamental issue. It does not arise from the existence of an abyss between us and others[2], but from our perplexing failure to ground meaning and understanding – even in our own case, in our own engagement with signs. We can never

[1] As I read them, the passages in the rule-following considerations on the teacher and the pupil make that point salient (see for example *PI* §208-213).

[2] As McDowell would put it, the regress is not an argument for scepticism about other minds. The point is not that there is an epistemological asymmetry between the first and third-person perspectives, such that I can know in my own case what the stop sign means to me, but can never know in the case of others. (see McDowell, 2000, 42)

secure a shared understanding of a sign, if the sign can never mean anything to anybody. The radical nature of the conclusion lies, one could say, in its complete generality.

To ground our own understanding of the stop sign, we are not bound to overt signs. For example, we might appeal to mental instructions we have previously given ourselves (rather than written ones), or to a chart of the various signposts we have memorised. But even then, what assurance would we have that we now understand these instructions or this chart in the same way we initially did? Is it not conceivable that such mental items could be variously interpreted? The mental chart can clearly be 'read' according to different schemas.[3] So how can it ever determine what the stop sign means to us? If the issue can be raised in the case of overt signs, like signposts or written instructions, it can equally be raised in the case of mental items. The regress threatens everything we ordinarily regard as having meaning – and so, more generally, as having intentional content. And surely, this includes what we have in mind.[4]

How then could the regress ever be stopped once it is accepted that understanding a sign (i.e. grasping its meaning) always hinges upon a mediating act of interpretation? For McDowell, the rules-as-rails mythology is an obvious attempt to put a halt to the regress. The mythology posits universals or supra-natural entities in a platonistic realm that rigidly and univocally mark what accords with them, prior to any engagement on our part. But it also insists that we can grasp such entities, in what would be analogous to "engaging mental wheels with […] objectively existing rails" (McDowell, 2000, 41). The stop sign, the written instructions, the mental chart might all be open to interpretation. These platonistic

[3] On this point, see Wittgenstein, *PI* §86.

[4] Or to phrase it differently: a sign is anything that we regard as having a meaning or as 'standing for' something else. Consequently, what we have in mind (in this case some instructions, or a chart) are signs – and so, are threatened by the regress.

entities however are not. And that is why our grasping them can ultimately ground meaning and understanding. As such, the mythology argues for the existence of a last interpretation that cannot itself be interpreted (see McDowell, 1984, 332).[5] However, McDowell contends that this only provides us with a false sense of security. Once we accept the initial appeal of the regress, the solidity of any stopping point will only ever be illusory – either because the assumed stopping point is not really a stopping point or because the assumed stopping point does not really exist (i.e. is mythological).[6] It might satisfy us for a time, but only for a time. The way out is to not embark on the regress at all, to show that it springs from a misconception. And according to McDowell, that's precisely what the rule-following considerations are designed to do. They are an attempt to rescue us from the philosophical bewitchment that makes the regress look unavoidable in the first place.

Now, I think this position, attributed to Wittgenstein, is also McDowell's own position (see especially McDowell, 1984, 1992, 2009 and 2013b). He too considers that the regress of interpretations can be avoided altogether if we reject the assumption that signs just 'stand there', that it is "only by virtue of being interpreted in one of various possible ways that they can impose a sorting of extra-mental items into those that accord with them and those that do not" (McDowell, 1992, 47).[7] While McDowell initially frames it in terms of "mental items", the point is clearly a general one.[8] His suggestion is not that, while overt signs (e.g. signposts, words, etc.) are normatively inert – and so incapable by

[5] As an illustration of this mythology, McDowell quotes a passage from the *Blue Book*: "What one wants to say is: "Every sign is capable of interpretation; but the meaning mustn't be capable of interpretation. It is the last interpretation." (See McDowell, 2013, 83-84 and *BB* 34).
[6] Whether or not McDowell is convincing on this point depends largely on whether or not the rules-as-rails mythology is indeed mythological. I think he is, though much more could be said on the matter.
[7] The phrase "stands there" is taken from Wittgenstein: "A rule stands there like a signpost" (*PI* §85). In this context, it gives expression to the misconception that signs are always in need of an interpretation (see McDowell, 1992, 47).
[8] See for example his discussion of the master thesis in McDowell 1992, 47.

themselves of sorting things into those that accord with them and those that do not –, thoughts at least are not. This would clearly commit him to the position he has already dismissed. It would amount to accepting the regress while claiming that it can be put to a halt. That is, it would imply something along these lines: of course, overt signs can always be variously interpreted. Things are not hopeless however: at one point, the interpretation will be a mental item and mental items are not themselves open to interpretation.[9]

This, surely, is not McDowell's position – nor is it Wittgenstein's. In a more recent paper[10], McDowell even explicitly maintains that "we must insist on not being deprived of the thought that a sign-post itself points the way, tells people which way to go" (McDowell, 2013b, 101). That is, he is adamant that it would be a mistake to think that, when we normally interact with them, signposts just stand there. Our understanding a signpost is, in general, an unmediated act; it is not dependent in any way on an interpretation. I say in general, because there can clearly be cases where the signpost remains silent to us (i.e. cases in which the signpost does not, by itself, tell us what to do or which way to go).[11] And in such cases, we might need further indications (i.e. we might need to be provided with an 'interpretation'). As I understand him, McDowell would not want to deny this. But he would insist that such cases are not representative of our interactions with signposts. The point is not that there are no circumstances under which the signpost telling us the way hinges on an interpretation. It is rather that, in ordinary circumstances, signposts in

[9] It is not to say, of course, that there are no interesting differences between overt and mental signs for McDowell.
[10] I have in mind "How Not to Read *Philosophical Investigations*: Brandom's Wittgenstein" (2013b, 96-111).
[11] This can be because we rarely encounter the signpost and do not remember what it means. This surely happens sometimes to apprentice drivers who are not yet well acquainted with the road signalisation – or to travellers in a foreign country. Or the signpost might fail to show us the direction because, it is somewhat ambiguous. Imagine for example if you were to encounter a sign with an upward arrow "↑" when it is possible either to go straight or to go up. (On the possibility of a sign being ambiguous, see also McDowell, 2009, 368).

themselves are not normatively inert.[12] And the same goes for words, whether spoken or written – and for any sign for that matter. That is what McDowell is suggesting in his more recent paper when he claims that the regress shows "that acting on an understanding cannot in general be acting on an interpretation of what is understood" (McDowell, 2013b, 103).[13]

1.3 Quietism and Anti-Reductionism

Now, McDowell often suggests that philosophy cannot contribute more to this particular region of thought. Its task is to show that the regress of interpretations is avoidable, that despite our initial anxiety the possibility of meaning and understanding was never really threatened. Once this has been done – once the regress has been shown to spring from a misconception about meaning and understanding – nothing more is needed. This, in a nutshell, is the quietist stance McDowell sees as central to the rule-following considerations and to Wittgenstein's later philosophy more generally (see for example McDowell, 2009 and McDowell 2013b). The thought that, for all signs, understanding is in general unmediated by interpretation does not call for constructive philosophy. It is not a somewhat vague insight about our dealing with signs that needs to be supplemented by a complete theory of meaning. That understanding is, in most circumstances, unmediated by interpretation is not a philosophical discovery that needs to be fleshed out, but something

[12] Compare to Wittgenstein's reply to the suggestion that one's understanding depends on one's having chosen a particular interpretation: "Not at all! A doubt was possible in certain circumstances. But that is not to say that I did doubt, or even could doubt" (*PI* §213)

[13] I think it is already McDowell's position in "Wittgenstein on Following a Rule", although it is not expressed as clearly then (see for example 1984, 338 where he says "But Wittgenstein's point is that this *dilemma* seems compulsory only on the assumption that understanding is always interpretation.". There he discusses the dilemma between the rule-as-rail mythology and the communitarian approach to rule-following, rather than the regress itself. The two are of course related, but it is not made explicit in that passage. Furthermore, McDowell does not really elaborate on that thought in this paper.)

that was always there in plain sight and of which we must remind ourselves when the temptation of the regress reappears.

Clearly McDowell considers this quietist stance to be a legitimate way of conceiving philosophy's task – in certain regions of thought at least – since it is a stance, he himself takes when discussing rule-following, the place of our conceptual capacity in nature, etc.[14] And this point is important. Many commentators recognise that Wittgenstein puts forward something along these lines. There are after all many textual evidences that seem to support it; Wittgenstein explicitly suggesting in various passages that philosophy is essentially therapeutic (*PI* §133 or §255), that its progress does not depend on any discovery or invention (*PI* §126), that its aim is to find peace, to free us from the illusions and misconceptions that lead to 'philosophical problems', and make them look pressing (*PI* §§110-111), etc.[15] However, there is also a widespread tendency among these commentators to decry this aspect of Wittgenstein's later philosophy – especially when it comes to rule-following. It is the case, for example, of Wright who considers that, in his later writings, Wittgenstein does actually raise important and pressing questions about the possibility of meaning and understanding, the nature of rule-following, etc. For him, Wittgenstein's stance is bound to be perplexing, if not disappointing: how can he bring these questions to the foreground and not even attempt to answer them? What could possibly justify his quietism in this situation?[16] It is also the case of Brandom who sees the regress of interpretations as calling for a new theory of norms and linguistic practices – a

[14] I have in mind, for example, the lectures 4 and 5 of *Mind and World* where he discusses his naturalised platonism (see McDowell 1996).
[15] In the same vein, although in more metaphorical terms, Wittgenstein also says that his aim in philosophy is to "show the fly the way out of the fly bottle" (*PI* §309).
[16] Wright's position is a bit more nuanced in recent papers (see for example Wright 2007), than it was previously (see his "Kripke's Account of the Argument Against Private Language" and some essays in his *Rails to Infinity*). But his disappointment towards Wittgenstein's quietism remains.

new theory however that Wittgenstein does not even attempt to provide on the ground of his quietism (see Brandom, 1998, 29)

As one can expect, McDowell is very critical of both Brandom's and Wright's take on the rule-following considerations (see McDowell, 1996, 175-176; 2009, 370 and 2013b, 104). But the disagreement is not merely exegetical. For Brandom and Wright, Wittgenstein hints at a response to the regress of interpretations in many passages, but he always falls short of satisfyingly fleshing it out. For example, he mentions the notions of custom and practice, but he does not elaborate on them nor does he precisely situate them in a systematic account of meaning and understanding. The issue is that, for Brandom and Wright, the initial moment of the regress is irresistible. Wittgenstein effectively shows us that signs, in themselves, just stand there – and so that invoking further signs (i.e. interpretations) in an attempt to ground meaning and understanding is hopeless.[17] He then rightly suggests that the regress can be stopped – and so that signs can still impose normative constraints on us –, if we can make sense of the idea that our engagement with signs is ultimately grounded on our participation in a practice. But he stops there, shying away at the task of accounting for that basic level that underpins meaning and understanding. Thus, for Brandom and Wright, the rule-following considerations do not offer a fully satisfying response to the regress of interpretations, remaining largely programmatic. Whereas for McDowell, everything that is needed to escape the threat of the regress is already there. So of course, the disagreement partly concerns how to understand what is achieved in the rule-following considerations. But this is more a matter

[17] See for example Brandom's comparing the regress of interpretations to Carroll's regress of inference. Brandom, 1998, 22).

of what counts as a response to the regress than a matter of what Wittgenstein is effectively trying to say.

Who is right then? Can one really be quietist about the possibility of meaning and understanding and the nature of rule-following – as McDowell claims one should be? Or are Brandom and Wright right to try to supplement Wittgenstein's remarks with a full-fledged – and so substantive – account of meaning and understanding? I want to suggest that the matter is a bit more complex than what is usually assumed. McDowell ultimately makes two distinct claims. As we have seen, he maintains that the only way to 'escape' the regress of interpretations is to not embark on the regress in the first place. And this is only possible, according to him, if we recognise that, *for any sign*, understanding is not in general a matter of interpretation. But this recognition in itself does not commit McDowell to quietism. After all, one could conceivably agree with McDowell's diagnosis and yet maintain that some kind of constructive philosophy is still called for.[18] That is to say, one could agree that rejecting the master thesis in the way McDowell suggests has to be part of a satisfying response to the regress, but contend that it ultimately needs to be supplemented – if not supported – by a *theory* or *substantive account* of meaning and understanding.

Now, we must avoid a misunderstanding. The master thesis consists in the thought that a sign can only ever impose a normative constraint on us under a particular interpretation (see McDowell, 1992, 45). And, as we have seen, McDowell forcefully rejects this thought. But it is important to realise that he is ultimately committed to more than the mere negation of the master thesis. Indeed, McDowell does not merely claim that

[18]The rejection of the master thesis might somewhat encourage the adoption of a quietist stance, but it does not, in itself, warrant it – despite what McDowell sometimes suggests. While rejecting the master thesis can be seen as a diagnostic – and therapeutic – move, it is not the only way to conceive it.

some signs are so constituted as to impose a normative constraint on us, without the need of any interpretation. He makes the stronger claim that no sign is, in itself, normatively inert – and so that all signs can, in themselves, sort things as according with them or going against them.[19] The underlying argument is that if one recognises the appeal of the regress in the first place – and so recognises that at least some signs are *intrinsically* normatively inert –, then it is very difficult to see how one can bring the regress to a halt without postulating something like interpretations that cannot themselves be interpreted.[20] And then the question of what is special about those signs that make them uninterpretable seems impossible to settle without accepting the 'rules-as-rails' mythology. Alternatively, one might dig deeper and say that the 'last interpretation' is not a queer entity but simply a kind of understanding that manifests itself in action.[21] However, this will not do for McDowell, partly because it 'undermentalizes' our engagement with signs (see McDowell, 1992, 50 and 2013a, 79-95).[22]

Thus, in rejecting the master thesis, McDowell ends up committing himself to the thought that there are not two kinds of signs corresponding to two kinds of understanding: one that is vulnerable to the regress of interpretations and another, more fundamental that is immune to it, one that gets the regress going and another that puts a halt to it.[23] But one

[19] As I have hinted at, I suspect he would also maintain that, for most signs, the vast majority of our everyday engagement with them are unmediated by interpretation. This is of course fully compatible with the recognition that, with any sign, our understanding can *sometimes* depend on an interpretation.
[20] See *supra* 1.2, footnote 5.
[21] With their insistence on the notion of practice, it is for example tempting to read Brandom and Wright as suggesting something along these lines.
[22] This is of course tightly related with McDowell's anti-reductionism about meaning and understanding. More on this later.
[23] Compare, for example, with Brandom's insistence that "norms that are explicit in the form of rules *presupposes* norms implicit in practice." (Brandom, 1998, 20). Of course, the regress would never end if norms were all explicit rules, but some norms are implicit in our practice and, as such, escape the threat of the regress… As I see it, Wright's most recent discussion of rule-following also repeats the same schema: of course, the *modus ponen* model of rule-following naturally leads to some form of regress if it were to cover all cases of rule-following, but it does not. There are also basic cases of rule-following that are structurally

could accept the same commitment without rejecting as irrelevant questions such as "what does understanding actually consist in, if not in an interpretation?", "what is it to grasp the meaning of a sign?", etc. As such, McDowell's claim that it can never be philosophy's job to answer such questions is essentially distinct from his rejection of the master thesis.[24]

We can now see how Brandom's and Wright's position on the issue is also twofold: contrary to McDowell, their rejection of the master thesis is less committal and they are willing to accept that the regress – or at least the initial moment of the regress – is unavoidable. For them, if the regress does not ultimately threaten meaning and understanding, it is not because it is illusory from the beginning, but because it has an end. Or to put it differently: while the regress is real for them, it is at least not infinite. However, this can – and must – be distinguished from their criticism of quietism. Indeed, Brandom and Wright do believe that that regress can only be satisfyingly answered by a proper theory

---

distinct. (see Wright, 2007, especially section III and IV). It is worth mentioning though that, in this paper, Wright does not appear to be as committed as he once was to the idea that these basic cases have to be understood in terms of our participation in a practice. Still, like Brandom, he seems to be suggesting that the distinction between an understanding mediated by an interpretation and an understanding unmediated by an interpretation is a *foundational* one. That is to say there seem to be for them two classes of signs: those whose understanding is mediated by an interpretation – and so depends on our grasping some sort of explicit rule – and those whose understanding is unmediated by an interpretation – and as such are best captured by Wright's basic cases of rule-following or by Brandom's implicit norms in practice.

If Wright is never that explicit, it is still rather telling that he consistently uses two different rules to illustrate the distinction between cases of rule-following that are assimilable to the *modus ponen* model and cases of rule-following that are basic (see Wright, 2001, especially his "Wittgenstein's Rule-Following Considerations and the Central Project of Theoretical Linguistics" and 2007).

It is worth noting that, in his criticism, McDowell does not explicitly accuse Brandom or Wright of positing "two kinds of signs corresponding to two kinds of understanding" – as we have done – but rather of positing "two correctness" (McDowell, 2013b, 109). However, I take these formulations to be ultimately equivalent.

[24] That being said, McDowell appears willing to concede that while philosophy cannot contribute more to our understanding of meaning and understanding, the sciences might. That's how I read for example a passage in *Mind and World* in which he says: "It would be one thing to give an evolutionary account of the fact that normal human maturation includes the acquisition of second nature, which involves responsiveness to meaning; it would be quite another thing to give a constitutive account of what responsiveness to meaning is" (McDowell, 1996, 124). His point seems to be that there might be an evolutionary account of the basic facts about our nature that enable us to understand signs (i.e. to grasp/be responsive to their meaning). But that it would be unreasonable to also expect a – philosophical – account of what this understanding amounts to.

of meaning and understanding.[25] But one can conceivably accept, as they do, that understanding a sign is in general mediated by an interpretation, except at a basic level and yet be satisfied with far less than a theory of what that basic level is. That is to say, one can conceivably accept a picture of signs that gets the regress going and yet maintain that the regress does not call for constructive philosophy. Maybe one could be satisfied by the mere recognition that there *must* be a basic level at which the regress stops. Surely this would amount to some sort of quietism, but it would not collapse into McDowell's position. For one thing, such a position would force us to accept that some signs (i.e. those that are not at the basic level) are normatively inert, which, as we have seen, McDowell forcefully rejects.

The question of whether or not understanding is really a matter of interpretation above an assumed basic level is not equivalent to the question of whether or not a quietist stance is warranted towards the regress. So, it might be possible, as I have suggested above, that McDowell is right concerning the necessity of freeing ourselves of the idea that signs – or at least some signs – are, in themselves, normatively inert, but wrong concerning the necessity of a philosophical account of our "responsiveness to meaning". It is, in any case, a suggestion I will take seriously in later chapters.

Now it is worth stressing that McDowell is also committed to some form of non-reductionism. According to him, it is simply not possible to reconstruct meaning, understanding – and other normative phenomena – out of materials that do not already presuppose these very phenomena. That is partly why he is so critical of Wright's initial response to the regress – and of communitarian accounts of meaning more generally. The

[25] Contrary to McDowell, they are also of the opinion that such a theory is a legitimate goal for philosophy.

problem with such accounts, according to him, is that they situate the basic level (i.e. the level at which the regress allegedly comes to a halt) so low that they end up unable to accommodate meaning and understanding at all. By only positing neutrally characterisable behavioural dispositions – and maybe accompanying feelings –, they deprive themselves of the necessary resources to make sense of "meaning's normative reach", of verbal behaviour being submitted to semantic constraints of right and wrong, etc.[26]

For McDowell, there is no way of making sense of verbal behaviour as endowed with – particular – meaning if we do not already characterise that behaviour in terms of its semantic content.[27] The fact that we mean *red* when we utter the word "red" is bound to remain out of our reach as long as we try to reconstruct it out of verbal dispositions characterised in semantically neutral terms. How then could we ever distinguish our meaning *red* by the word "red" – i.e. our being committed to using "red" according to that meaning – from the mere regular production of the sound "red" in these and those circumstances? For McDowell, the only way to avoid these difficulties is to resist the temptation to reduce our verbal behaviour to a series of neutrally characterisable dispositions in the first place. We do not merely have a disposition to produce the sound "red" in these and those circumstances – like a cleverly designed automaton or a properly

[26] The phrase "meaning's normative reach" is from McDowell, 1992, 48. The idea is not merely that verbal behaviour is submitted to normative constraints, but that these constraints are essentially semantic. For a general discussion of the issue surrounding communitarian account of meaning, see McDowell, 1984, 340-341.

[27] McDowell often cites the following passage in Wittgenstein's later writings to support his non-reductionism about meaning and understanding: "Following according to the rule is FUNDAMENTAL to our language game" (*RMF-VI* §28. see also McDowell, 1984, 340-341 and 2013b, 102). However, it is not clear to me that this passage really supports McDowell's position, especially if we understand rule-following, as he does, as simply "a gloss on the idea of acting in accordance with a conception of correctness" (McDowell, 2013b, 102). If rule-following really just stands for any kind of normative constraints, then one could conceivably claim that bedrock – or the basic level – is 'populated' by normative constraints, but not by meaning and understanding. That is to say, while the passage from Wittgenstein could support an anti-reductionist account of meaning and understanding, it could equally support a partially reductionist account, like the one recently put forward by Ginsborg (see Ginsborg, 2011b and 2012).

trained parrot might have –, but also a disposition to mean *red* by "red".[28] The fact that signs are endowed with meaning – and that we respond to them with understanding – has to be recognised as self-standing: it cannot be grounded on more basic facts (i.e. facts that would be characterisable in non-semantic terms). To be sure, the point is a general one for McDowell. It concerns our overt verbal behaviour – and overt signs more generally – as much as it concerns our thoughts. The fact, for example, that I mean *this is red* when I point at a red object and utter "red" is only available to us if we are willing to describe that verbal behaviour as an assertion that *this is red*. In the same way, the fact that I think *this is red* can only be secured if we are allowed to describe my mental event as a case of thinking that *this is red*.[29]

In light of these considerations, we can start to see how McDowell's anti-reductionism about meaning and understanding is distinct from his quietism. His anti-reductionism commits him to the metaphysical claim that facts about meaning and understanding are *bedrock* – and so are not reducible to more basic facts about our verbal behaviour, our thoughts, etc. His quietist stance, on the other hand, is not so much a positive claim – metaphysical or otherwise –- as it is a methodological commitment. It does not stipulate anything about meaning and understanding nor about what can be explained about them; simply that no such explanation is ever *required*.

Now, obviously if McDowell's adoption of a quietist stance did imply some positive claim on what *can* be explained – and so on what *can* be known and said –about meaning and understanding, then it might give credence to his anti-reductionism, or vice

[28] On this point, see McDowell, 1984, 340-342.
[29] On this point, see McDowell's discussion of the thought that people are talking about me in the next room. (see McDowell, 1992, 46-48).

versa. After all, it would appear contradictory to claim that an explanation is possible and yet, maintain that there is, as it were, nothing to be explained. (How could it be possible to dig deeper if we are already at bedrock?). Relatedly, it would appear unreasonable to deny the possibility of any explanation, while maintaining that there is actually something to be explained. (If we have not yet reached bedrock, then how could it be impossible, in principle, to dig deeper?) But this proves to be a mistaken view of what is entailed by one's adoption of a quietist stance: there is ultimately no such relation of implication between such a commitment and any form of anti-reductionism.[30]

Why does this matter? Well, if our project is to lay the groundwork for an account of rule-following that is both substantive and partially reductive, then we will need to depart from McDowell on these matters. And we can only do so, if we are clear as to what these commitments (to a quietist stance and to a form of anti-reductionism) respectively entail – and how they relate, or fail to relate, to one another.

When it comes to his anti-reductionism, McDowell appears to be largely motivated by the alleged failure of Wright's and other communitarian accounts to ground meaning and understanding on non-semantic and non-normative facts about our behaviour. McDowell might be right that these accounts are unsuccessful – and that they are unsuccessful *because* they dig too low. But this does not, in itself, prove that facts about meaning and understanding, as he conceives of them, are bedrock. For one thing, it does not rule out the possibility that there might be a level below meaning and understanding, but above brute facts about our behaviour. Such a level might be hospitable to a more basic

[30] We'll say more about this in subsequent sections (see infra 1.5-1.7).

kind of intentionality – or at least of normativity – that is bedrock and from which we can ultimately explain meaning and understanding…[31]

This will be the starting point of our attempt to flesh out a partially reductive view of meaning and understanding – and the cornerstone of our opposition to McDowell on that particular issue. In contrast, our attempt to put forward a substantive account of rule-following – or at least to lay the groundwork for such an account – will need to rest on other grounds and will ultimately mark a distinct opposition to McDowell. On the issue of locating bedrock, our disagreement with McDowell will be metaphysical. On the issue of constructive philosophy versus quietism, our disagreement will rather be methodological.

1.4 Objectivity and Guidance

So far, I have suggested that we can be critical of McDowell's anti-reductionism and quietism, while preserving his insight that signs are precisely the kind of things that can be *intrinsically* meaningful and so, in themselves, impose normative constraints on us. To conclude, I would like to briefly consider how McDowell situates – and further fleshes out – this insight in relation to, what I've called, the objectivity and guidance conditions on rule-following.[32] As we have suggested, the objectivity condition specifies that a satisfying account of rule-following must accommodate the insight that the requirements of a rule precede or, at any rate, are independent of our attempt to know them. Of course, it is not to deny that our following a rule might ultimately depend on our being successful in this

[31] Hannah Ginsborg proposes one way to flesh out the possibility of a basic level of normativity, more fundamental than any form of intentionality (see Ginsborg, 2011a, 2011b, 2012, 2020 and 2021), In contrast, I will propose in chapter 4 a more basic level of intentionality to the one envisioned by McDowell and from which a *sui generis* kind of normativity emerges.

[32] Since McDowell considers the expression "following a rule" to be "a gloss on the idea of acting in the light of a conception of correctness" (McDowell, 2013b, 99), it clearly encompasses the idea of a sign having meaning and so imposing specific normative constraints on us.

endeavour. The point is rather that the requirements of the rule cannot collapse with what we take them to be. What, as a matter of fact, accords with the expression of a rule cannot simply be what we think accords with such an expression. Otherwise, we could not talk of anything being in accord with the expression of a rule – or going against it. If we frame this in terms of meaning rather than in terms of rule or of requirements of a rule, we get something like this: whether or not something accords with a sign, on the ground of its meaning, has to be independent of our alleged awareness or appreciation of its meaning.[33]

As for the guidance condition, it implies, in part, that the requirements of the rule must be in principle available to us, if following a rule is to be possible at all. The rule has to be such that it is within reach of our cognitive capacities. The requirements of the rule might not be up to us, but it must be possible for us to be receptive and ultimately responsive to them.[34] There is, however more to the condition: the requirements of the rule must be available to us, but they also must be available in the right way. Not only must we be aware of the requirements of a rule upon acting, for our action to be an instance of rule-following. Our awareness of the requirements also needs to somehow inform our acting the way we do.

Now, McDowell agrees that a satisfying account of rule-following must preserve a somewhat familiar and intuitive notion of objectivity, which is to say that he recognises that there is indeed an objectivity condition on rule-following (see McDowell, 1984, 325-

[33] A point of clarification is in order, we talk of "a rule", "the expression of a rule" and "the requirements of a rule", but it might not be obviously clear how these expressions relate to one another. I take that the expression of a rule – if it does not just stand there – is a meaningful sign (i.e. a sign endowed with meaning). The rule *per se* and its requirements would rather correspond to the meaning of the sign – and its related implications for what accords, or fails to accord, with the sign.

[34] Or to put it differently, while we can fail in our attempt to know these requirements, things must be such that we can sometimes succeed.

326 and 334).[35] And he considers that his stance on meaning and understanding actually satisfies that condition. However, in his canonical papers on the topic, he ultimately says very little about how to precisely flesh this out.[36] Luckily, McDowell is a lot more explicit when he discusses the objectivity of aesthetic values or of the demands of virtue.[37] And since he sees fundamental similarities between these particular regions of reality and the realm of meaning and understanding, it is not unreasonable to assume that the conclusions he draws about the former can shed light on his stance on the latter.

For McDowell, the crucial move is to acknowledge that things can be objectively there – in the sense of being there *as they are* independently of any particular judgment or set of judgments we make about them[38] – and yet depend ontologically on our receptivity, or sensitivity, to them. Take for example secondary qualities, like colours. McDowell wants to say both that colours are genuinely part of the "fabric of the world" and that they are not conceivable independently of beings able to perceive and respond to them (see McDowell, 1983, 5). When I see an object as being red, I am sensitive to a property genuinely possessed by that object. It is not just that the object appears to be red – or causes in me the *impression* of red: the object *is* red (see McDowell, 2011, 115). And for McDowell, this is perfectly compatible with the recognition that things can be red only because things can be experienced as red.

---

[35] He sometimes expresses this thought by saying that "ratification-independence" is constitutive of our idea of meaning (see McDowell, 1992, 48).

[36] The papers in question are of course "Wittgenstein on Following a Rule" and "Meaning and Intentionality in Wittgenstein's Later Philosophy" (see McDowell 1984 and 1992).

[37] On the question of aesthetic values, see McDowell 1983. The question of the demands of virtue is the focus of a few papers (see McDowell, 1978, 1979, 2000). One can also find a few passages on values more generally (see McDowell, 1996, 2011).

[38] Insisting on the independence from any particular *set* of judgments enables us to preserve the insight that we can be collectively mistaken. How things are cannot collapse with how I think things are. But neither can it collapse with how *we* think things are (On the importance of preserving this insight, see Blackburn, 2005, 171 and 173 and Wright 2007).

For McDowell, the same goes for the things that populate the space of reason – whether it is values (including aesthetic values), virtues or ultimately meanings. As he puts it: these things "are not brutely there – not there independently of our sensibility – any more than colours are: though, as with colours, this does not stop us supposing that they are there independently of any particular apparent experience of them" (McDowell, 2011, 120). The two claims only appear to be in tension because we are tempted by a misguided picture of what it is to be objectively there. Of course, if we maintain that only things that are conceivable in complete isolation from any particular perspective or point of view are genuinely part of the fabric of the world, then there does not seem to be place for colours, any more than there is for values or meanings. They are, at best, appearances that need to be explained away (on this point, see McDowell, 1983). But, as McDowell forcefully insists, we should not accept that what is real is what lies there before the involvement of any sensibility. That would amount to erasing whole regions of reality.[39]

One might feel that there is an important disanalogy between secondary qualities and the things that populate the space of reason: whereas our sensibility to colours, etc. appears to be wholly dependent on our biological constitution, our sensibility to values and meanings appears to be – at least partly – shaped by our initiation into the relevant practice. I can see red without any training, or so the thought goes, but I can only 'see' the meaning of your word or the value of your action given a certain upbringing. However, it should be clear that this is not McDowell's position. For him, there is no disanalogy because our

[39] I assume McDowell is somewhat committed to the thought that we simply cannot reconstruct colours out of materials that are 'brutely' there.

sensibility to colours turns out to be dependent on conceptual capacities that also demands a certain upbringing to be actualised (see McDowell, 1996).[40]

For McDowell, the fact that things are *inconceivable* in complete abstraction from our sensibility – and even responsiveness –- to them does not speak against them being part of the fabric of the world. Take here virtues: for McDowell, virtues – or at any rate a conception of ethics revolving around the possession of virtues – involves "requirements of reason that are there whether we know it or not […]" (McDowell, 1996, 79). That is, virtues involve requirements of reason that are objectively there. At the same time, it is clear that such requirements are only conceivable within the context of an ethical life – and so, from the perspective of an agent who is appropriately sensitive and ultimately responsive to them. If we cannot conceive what it is to be red in complete isolation from what it is to appear red, neither can we conceive of a virtue – for example courage – in complete isolation from the givenness, to an agent, of what courage requires.

This naturally leads us to the question of guidance. Here, there is however a genuine and important disanalogy between secondary qualities and things that populate the space of reasons. Take colours: they obviously are epistemically available to us – except, maybe if we are colour blind. And their availability is rather straightforward. We can know that an object is red simply on the ground of our having the relevant perceptual capacities– which, for McDowell, involves the constitution of our sense organs as much as our conceptual repertoire.

But there is no obvious sense in which secondary qualities *constitutively* guide how we respond to – or engage with – them. Of course, one might recognise a signpost – a stop

[40] The point is that, for McDowell, something cannot look red to me – look to have the property of being red – without initiation into a practice.

sign, for example – in part because of its colour. And so, if one stops at the stop sign, there might be a sense in which one's action has been 'informed' by one's perception – and so awareness – of the colour of the stop sign. But clearly the stop sign does not tell one what to do *on the ground* of its colour. Or to put it differently: to appear a certain colour (e.g. red), in this case or in general, does not imply any sort of requirements on how one should respond.

When it comes to the guidance condition, the case of virtues proves more informative. McDowell repeatedly insists, not merely that virtues and their requirements are an objective matter, but also that, in being initiated into the relevant practices "our eyes are opened to them" (see McDowell, 1996, 79). Virtues and their requirements are epistemically available; they are within our cognitive reach. Here, I think, McDowell should be taken quite literally: in saying that our eyes are – or at least can be – opened to virtues and their requirements, McDowell is not merely offering a metaphor of them being available to us: he is really making a claim about *how* they are available to us.. He is rather explicit about it in an earlier paper: one's responsiveness to virtues and their requirements depends on one's sensibility to these requirements. And that sensibility really is "a sort of perceptual capacity" (McDowell, 1979, 332).

Importantly, McDowell goes further: such a perceptual capacity delivers a form of knowledge or understanding of virtues and their requirements. To *see*, for example, what courage requires in a given situation *is* to know or understand what is required of oneself in that situation. It *tells* one what to do and so *guides* one's response to the situation. (see McDowell 1979, as well as McDowell, 1978 and 2000.)

What about meanings? In his canonical papers on rule-following, McDowell does talk a few times of seeing the meaning of a signpost or of hearing the meaning in someone else's utterances (see McDowell, 1984, 351 and 355). In later writings, he persists and relates our understanding of a sign (i.e. our responsiveness to its meaning) to our *sentient* capacities (see McDowell, 1996, 77). Now, some might be tempted to see these passages as largely metaphorical, but I think that would be a mistake. In my opinion – and in light of what he says about virtues –, McDowell is really suggesting that meanings can be given to us in perceptual – or quasi-perceptual – experiences. Ultimately, and just like with virtues, I think he also ends up suggesting that, to grasp the meaning of a sign – to understand what accords with that sign – need not amount to anything more than to perceive its meaning.[41] In the same way that, for McDowell, we have the relevant perceptual capacities to be sensitive to colours or virtues, we have the relevant perceptual capacities to be sensitive to meanings. But unlike secondary qualities, though very much like virtues, our sensitivity to meaningful signs amount to a form of knowledge – or understanding – of certain requirements (i.e. of what accords or fails to accord with them). Consequently, there is, in general, no gap between our perception of meaningful signs and our understanding of them. No gap for which further cognitive steps would be required.

That, I think, is central to McDowell's position on rule-following. If we can resist the thought that signs, in themselves, are necessarily normatively inert – that is, if we can make room for the thought that signs can be *intrinsically* meaningful –, it is because we can make sense of the meaning of a sign being an objective matter, in the way just sketched.

[41] We can thus better understand McDowell's insistence that understanding should be considered a "definite mental state" (see McDowell, 2013a).

And if this is helpful in alleviating the threat of the regress, it is because we can make sense of such a meaning being (perceptually) given to us in a direct and unmediated way.

Take McDowell's recent suggestion that, in the case of the signpost, there need not be "an intellectual move that mediates between seeing the signpost and knowing which way it tells one to go" (McDowell, 2009, 367). What is implicitly assumed here is that, in our everyday life, we do not merely engage with physical objects, marks or sounds, but with signs which are intrinsically meaningful. And we can do so, because we have the relevant perceptual capacities. Once this is acknowledged, then the temptation that gets the regress going – i.e. the thought that there *must* be further "intellectual moves", "cognitive steps" or "interpretations" mediating between our perception of the *meaningful* signpost and our knowledge of which way to go – simply disappears.

1.5 Making Philosophical Claims

So far, we have argued that McDowell takes seriously Wittgenstein's purported quietism and himself adopts such a stance towards the regress of interpretations. But one might feel that, in discussing these questions, we have fallen short of really explaining what it is to be quietist – that we have merely repeated slogans (e.g. of philosophy as therapy, of description versus explanation, etc.) that are ultimately not clearer than the philosophical stance they are intended to illuminate. And it seems difficult to dismiss this worry outright as simply uncharitable. Of course, there is no denying that McDowell takes his response to the regress of interpretations to be a quietist one – and, as such, to be very much in line with Wittgenstein's (see most notably McDowell 2009). However, it is not always obvious why we should regard it as such. As we have seen, it is partly on exegetical grounds that

he is critical of Wright and Brandom. For him, their failure to take seriously Wittgenstein's quietism contribute to their misunderstanding of what the rule-following considerations actually achieved – hence his thinly veiled accusation that, in trying to make the rule-following considerations fit their respective philosophical project, they are ultimately distorting them (see for example McDowell 1984, 2009 and 2013b). But, his criticism of Wright and Brandom also involves a disagreement about the nature of our engagement with signs and of our sensibility to meanings – disagreement which is distinct from and independent of the aforementioned exegetical dispute. After all, the question of our relationship to signs can be addressed without relying on what Wittgenstein has to say on the matter – even if, to some extent, we are still indebted to him for the way the question is framed. Brandom and Wright are certainly not ignoring Wittgenstein's remarks on rule-following, nor are they dismissive of them. Nonetheless, what they are putting forward is not merely a commentary on these remarks, but a tentative account of the phenomenon these remarks are circumscribing. And so, it cannot be judged or dismissed only on exegetical grounds.

Obviously, this is not lost on McDowell; I am not trying to suggest that his criticism somewhat conflates these various points of disagreement. On the contrary, my suggestion is that, in his responses to Brandom and Wright, McDowell purposely engages, not only with their respective readings of the rule-following considerations, but also with what they have to say about the phenomenon of rule-following itself.[42] And it is in this latter engagement that the issue of quietism comes most clearly into focus. In answering the

[42] This does not mean, of course, that he does sometimes restrict his criticism to exegetical matters. (See, most notably, McDowell, 2013b. The very title of that essay, "How Not to Read *Philosophical Investigations*: Brandom's Wittgenstein", is rather telling.)

regress of interpretations, Brandom and Wright say something positive about the nature of our relationship to signs and of our responsiveness to meanings. They do not merely highlight what must be true of our involvement with signs, if the regress of interpretations is to not be endless; they also make claims about what is, as a matter of fact, true of this involvement. In a word, they propose to explain – if not to offer a substantive account of – what it is to follow a rule; not just to put a halt to the regress. (The point here is that a substantive account of rule-following is more than just a response to the regress. So, while a substantive account of rule-following must certainly show how the regress is not infinite, this cannot be the only criterion of whether or not such an account is satisfying.) But then, where does this leave McDowell? As we have argued, McDowell too says something positive about the nature of our relationship to signs and of our responsiveness to meanings. However, and contrary to Brandom and Wright, he does not consider himself to be providing an account of rule-following in doing so; merely to be dissolving the threat of the regress. This is of course not surprising and a very clear expression of his quietist stance.

Still, it is important to note that, on the question of our involvement with signs, McDowell's criticism of Brandom and Wright does not appear to primarily depend on such considerations. It is not to say of course that his quietism does not, in some way, tint his discussion of their respective account, but it is certainly not the focal point. Indeed, McDowell does not seem to be dismissing what Brandom and Wright have to say on rule-following on the ground that they misjudge the nature of the philosophical task imposed by the regress. Of course, McDowell does think that they are misguided in thinking that the regress calls for an account of rule-following, but ultimately that's not why he considers

their respective account to be wrong. As we have seen, McDowell raises three kinds of concerns against their accounts: (i) in digging below the level of meaning and understanding, they deprive themselves of the material necessary to explain meaning and understanding (see especially McDowell 1984), (ii) they 'undermentalise' our involvement with signs (see McDowell 1992) and (iii) in accepting the initial moment of the regress of interpretations, they remain committed to the assumption that makes the regress inescapable (see McDowell 2009 and 2013b).

Noticeably, McDowell frames his own proposal as a way to respond to these concerns. If it is recognised that signs in general are not in themselves normatively inert, then the regress can be avoided altogether – and not simply pushed back. In the same way, if it is accepted that we perceive the meaning of signs, then the problems both of 'digging below bedrock' and of 'undermentalising' our grasp and so our understanding of said meanings disappear. As such, there seems to be a sense in which McDowell is not simply hinting at how to avoid the regress, but is also arguing for a certain way of conceiving our involvement with signs. This is especially salient in the thought, we have attributed to him, that the meaning of signs can be – and often are – given to us in perception. For McDowell, all that is needed to avoid the regress of interpretations is to recognise that signs in general are not normatively inert. So why does he also make claims about the nature of our responsiveness to meaning – even going so far as to propose a story of how we come to manifest such a responsiveness? If McDowell's proposal was restricted to the modest claim that, in general, signs in themselves are not normatively inert, then there might be more plausibility in characterising it as simply an attempt to free ourselves from a problematic assumption. However, and as we have suggested in the previous section, his proposal

seems rather to amount to this: if, in general, signs in themselves are not normatively inert, it is because of something intrinsic to them (i.e. their meaning). What accord or fails to accord with them is not just an objective matter, it is an objective matter grounded on their meaning. And so, if signs, in themselves, can *tell* us what accords or fails to accord with them, it is, in part because they are the kind of things that they are (i.e. intrinsically meaningful). But, ultimately, it is also because we are sensitive to their intrinsic meaningfulness. And our sensibility here is really a perceptual capacity that is shaped by our initiation into the appropriate practices and which delivers a form of knowledge or understanding. To conceive signs, in themselves, as being necessarily silent is to mistakenly conceive them, and ourselves, in complete abstraction from a particular upbringing, or *Bildung*.

Evidently, this says something positive about our involvement with signs. It even looks like an explanation or at least the sketch of an explanation: there is a reason why signs are not, in themselves, normatively inert and it is not simply that the regress would otherwise be inevitable. Of course, the threat of the regress makes such a conclusion appealing, but that is not what ultimately grounds it. For all its quietist gloss, McDowell's proposal does appear to offer a reasoned and structured elucidation of what is actually going on when we understand a sign. That is to say: in criticising Brandom's and Wright's account, McDowell also seems to be putting forward substantive philosophical claims. But then, if we accept this characterisation – or at the very least the gist of it –, it does look as if McDowell's position is unstable. He is claiming to adopt a quietist stance towards the regress of interpretations, but – in apparent contradiction with this very claim – he is also making substantive claims about how we actually interact with signs and grasp meanings.

While McDowell insists that his response to the regress of interpretation is not so much the solution to a philosophical problem, but rather the dissolution of a misunderstanding, one can be perplexed as to what the difference really amounts to.

To be sure, one finds the same tension in Wittgenstein's discussion of rule-following – and in his later philosophy more generally. For example, one often gets the sense that Wittgenstein makes positive claims about what is actually involved in our following certain rules (e.g. the rule '+2') or in our engaging with signs of various kinds. Likewise, his alleged method – or rather, methods (*PI* §133) – appear to produce philosophical arguments to which constructive philosophy is far from inhospitable.

And so, the question appears unavoidable: independently of what they claim to be doing, are McDowell and Wittgenstein really doing something different? One underlying worry is that McDowell and Wittgenstein might ultimately fail to conform to their own methodological commitments: they repeatedly claim to be quietist, but, in practice, it might turn out that they consistently disregard the implication of adopting such a stance. That is to say: there might be a mismatch between how they think philosophy should be done and how they actually do philosophy. But the question reveals another – potentially more damaging – worry: our unease might actually concern the very intelligibility of what it is to be quietist. The adoption of such a stance is often conceived as a rejection of – and alternative to – more orthodox ways of engaging with philosophical difficulties. But then, what if the distinctiveness of the quietist stance is only superficial and does not withstand closer scrutiny? Again, McDowell and Wittgenstein would not really be doing something different. However, in this case, it would not be because they are guilty of methodological dissonance: adopting a quietist stance would simply turn out to be an empty commitment.

There is yet another possibility: the failure might not lie with McDowell and Wittgenstein in particular, nor with the intelligibility of quietism *per se*. It might rather stem from a broader misunderstanding of what orthodox approaches to philosophy provide – or, at least, claim to provide. That is, there might be a tendency to overestimate the implications of making substantive philosophical claims and ultimately to inflate what philosophical explanations aim to achieve. (On this point, the fault might be as much on those who consider themselves to be providing such explanations as it is on those, like McDowell and Wittgenstein, who are critical of them.) As such, McDowell and Wittgenstein might be responding to a distorted picture of what it is to put forward a philosophical account of a phenomenon, to explain it philosophically. Quietism might be perfectly intelligible; simply it might be proposed as an alternative to a stance that is never really adopted – and potentially never even really defended. McDowell and Wittgenstein would then not be guilty of methodological dissonance, but it would not necessarily follow that they are really doing something different. So, even in this case, adopting a quietist stance as an alternative to more orthodox approaches to philosophy might still turn out to be an empty commitment.[43]

Our interest in this question is twofold. Firstly, it is a point of intellectual rigor. So far, we have distinguished three strands in McDowell's discussion of rule-following: a commitment to quietism, a form of anti-reductionism about meaning and understanding,

[43] If we can raise the question of the intelligibility of adopting a quietist stance, we can raise the same question for more orthodox approaches to philosophy. So then, there might be two distinct versions of the thought that quietism might emerge as a response to a philosophical stance that is never really adopted. It might be that the orthodox approaches to philosophy are intelligible, but are, *as a matter of fact*, never adopted. Alternatively, it might be that these approaches are not even intelligible to begin with and that is *why* they are never *really* adopted… In this latter case, though, people might still *be under the illusion* that they are doing philosophy in that way – and so might still defend such approaches. (Wittgenstein's own claim that the method(s) of *Philosophical Investigations* destroy nothing but "houses of cards" might be one way to flesh out that thought (*PI* §118).)

and a number of claims about our involvement with signs. We want to make the case that it is, in principle, possible to agree with McDowell on that third strand of his view (i.e. the claims about our involvement with signs) without agreeing on – or at least accepting – the other two. So, obviously: if it turns out that McDowell's commitment to quietism is incompatible with his claims about our relationships to signs and our responsiveness to meanings, then we would have the assurance that, in abstracting them from such a way of doing philosophy, we are ultimately not distorting them. The same would of course be true if quietism turns out to be nothing more than empty slogans – either because it is in itself unintelligible or because it is not fundamentally different from more orthodox ways of doing philosophy. In the former case, we would be satisfied that quietism can be abandoned altogether; in the latter, that there is, in principle, no difficulty in situating McDowell's claims about our engagement with signs and our responsiveness to meanings within the context of a more orthodox approach to philosophy. Secondly, it offers us the opportunity to reflect on our own methodological commitment. The underlying aim of our project is to lay down the groundwork for a substantive account of rule-following. So, even if these worries prove to be unfounded, the question of what distinguishes a philosophically substantive account of a phenomenon from a quietist response to a philosophical difficulty remains of central importance. Indeed, in making these worries explicit, it becomes apparent that our unclarity is not merely about quietism, but equally about the nature of philosophical explanation – and so about what it is to put forward a philosophically substantive account, generally. It reveals a blind spot in our self-understanding that warrants further reflection.

1.6 Philosophy without End

I can only offer a few remarks on quietism – remarks which will certainly not settle the matter once and for all, but which will hopefully suggest a way forward. One issue is the precise relationship between its various expressions. We have so far mostly framed quietism as a distrust of philosophical explanation and as a refusal to make substantive philosophical claims – though we have also acknowledged that some of its expressions more explicitly encourage us to conceive philosophy in therapeutic terms.[44] Still, it is not clear how both these ideas are related. For example, is it possible to defend a broadly therapeutic conception of philosophy, while purposely making substantive philosophical claims? Conversely, could one refuse to engage in constructive philosophy and yet deny that philosophy's aim is essentially therapeutic? Another issue, of course, is the extent to which revealing the roots of our misunderstandings and ultimately dissolving philosophical problems is really different from putting forward theses and defending theories. If, as we have suggested, McDowell's quietism like Wittgenstein's does not amount to a refusal to make positive claims or to offer arguments in favour of those claims, then the distinctiveness of the quietist stance, if there is any, must rest on other grounds.

Now, there's one methodological claim repeatedly made by Wittgenstein that is often overlooked – or at least whose importance is not always fully appreciated. And I want to suggest that it is key to appreciate his broader commitment to quietism. Already in the preface of *Philosophical Investigations*, Wittgenstein hints at the *open-ended* nature of his investigations. (The contrast with the *Tractatus*'s preface, in which he boldly claims to

[44] There have been some attempts to contrast a liberatory reading of *Philosophical Investigations*' professed method with a therapeutic reading. (see most recently Read 2021). I am not entirely convinced by the move, but even if there was indeed a genuine contrast between both readings, I do not think it would affect the line of argument we are pursuing here.

have solved all philosophical problems once and for all, is striking.) The point is not so much that *Philosophical Investigations* have been left unfinished, that they are an incomplete work, but rather that their subject-matter resists – or at least is conceived as resisting – a complete and definitive elucidation. Of course, Wittgenstein does express a certain dissatisfaction towards *Philosophical Investigations* – and towards their various failures of expression. For example, he talks of the process of writing *Investigations*, as one of throwing away badly written remarks, of salvaging a number of "half-decent ones" by rearranging them and cutting them down, etc. He also ends the preface by saying: "I should have liked to produce a good book. It has not turned out that way, but the time is past in which I could improve it" (*PI* preface). However, this should not overshadow his *insistence* that the form they end up taking has ultimately to do with the very nature of what is investigated. That *Philosophical Investigations* approaches philosophical questions from "different directions", that they constantly come back to topics they have already explored, that they refuse to impose them a definite order, all these features are not in themselves failures of expression. They are rather the manifestations of a methodological commitment: the refusal to take philosophy as something that can be done once and for all. Wittgenstein makes essentially the same point in *Zettel*:

> Disquiet in philosophy might be said to arise from looking at philosophy wrongly, seeing it wrong, namely as if it were divided into (infinite) longitudinal strips instead of into (finite) cross strips. This inversion in our conception produces the *greatest* difficulty. So we try as it were to grasp the unlimited strips and complain that it cannot be done piecemeal. To be sure it cannot, if by a piece one means an infinite longitudinal strip. But it may well be done, if one means a cross-strip.—But in that case we never get to the end of our work!—Of course not, for it has no end. (*Z* §447)

Of course, if solving – or dissolving – philosophical difficulties is a matter of grasping some infinite longitudinal strips, then Wittgenstein's proposed method is bound to appear dissatisfying. Indeed, there would then be a fundamental mismatch between the form of his investigations and the nature of what is investigated. But, Wittgenstein suggests, our engagement with philosophical difficulties is rather akin to grasping finite cross strips cutting through infinite longitudinal strips. Consequently, there is no problem in principle to engage with philosophical difficulties as they appear and come into focus. A solution for all of them does not need to be available at once. The mystery of how infinite longitudinal strips could ever be within our reach in their entirety is shown to be irrelevant, of no real concern for philosophy. However, a concession is necessary: if philosophy can be done piecemeal in that fashion, it is at the condition of accepting that its results are always only partial.

We must be careful here. The picture of philosophical difficulties as finite cross strips rather than infinite longitudinal strips can quickly become misleading. The picture is helpful in showing how our engagement with philosophy can be fruitful even if it does not yield a complete elucidation of all difficulties at once. But it also has the flaw of suggesting that philosophical difficulties stand in complete isolation from one another. One might thus be tempted to conclude that, taken individually, they might be addressed and solved definitely. Philosophy's task would be infinite, but at least its progress – however modest – would be unshakeable. To be sure, this is not Wittgenstein's position – he does not endorse such claims, nor does he do philosophy in a way that assumes such an endorsement. To begin with, it clearly goes against his insistence that philosophical difficulties arise from a deficit in surveyability, from our failure to get an overview of our

language-games. For him, understanding in philosophy is a matter of "seeing connections" of arriving at a "surveyable representation". And of course, there is no way to circumscribe in advance the scope and diversity of these connections. Their discovery – Wittgenstein also talks of "invention" – is always contingent on our actual engagement with particular cases (see *PI* §122). Since new language-games can always emerge – and with them new connections –, we never have the assurance that the philosophical difficulties we consider solved will not require our attention again at some later point.[45]

There is a similar danger to misread Wittgenstein's assertion that, in engaging in philosophy, we seek complete clarity (see *PI* §133). The remark is often read as saying something like this: of course, such a complete clarity has proven elusive in practice, but it can be attained in principle and so, can reasonably be our standard of success. On this reading, Wittgenstein would be of the opinion that philosophy that fails to reach complete clarity – even on particular philosophical questions – is at best deficient, at worst defective. And evidently, this would fit uneasily with the humble profession made in the preface of *Philosophical Investigations*. The passages would need to be somewhat 'explained' away. But I think this is all wrong: Wittgenstein is not endorsing the idea that philosophy's only possible success is complete clarity; he is criticising it. As I see it, he is ultimately warning us off the temptation to think that philosophy would only be worth doing if such a complete clarity was to be reached. Philosophy might find its peace in complete clarity (see *PI* §133),

[45] On this point, Wittgenstein offers a rather striking analogy: our language is like "an ancient city: a maze of little streets and squares, of old and new houses […] and all this surrounded by a multitude of new suburbs with straight and regular streets and uniform houses" (*PI* §18). Baker and Hacker make no mention of it in their analytical commentary, but it is difficult to not take this as a nod to a famous passage of the *Meditations*, in which Descartes offers a similar analogy (see Baker & Hacker, 2005, 71-72 and Descartes, 1960). However, whereas Descartes is adamant that the most accomplished cities – like the most accomplished philosophies – really are those that are built at once, from scratch, Wittgenstein rather warns us against such a temptation, insisting that we must not "interfere in any way" with language (*PI* §124).

but that only means that philosophy can never find its peace. (Less provocatively, we could say that, for Wittgenstein, complete clarity is the asymptote towards which philosophy tends and, hopefully, progressively gets closer, without however ever reaching it. Indeed, my suggestion is not that for him, we should not aim for complete clarity, but rather that our inability to reach it should not be seen as a failure.)

And of course, this fits nicely with the idea that philosophy is a series of therapies – rather, let's say, than a cure. We can solve – or dissolve – philosophical difficulties, but we can only do it piecemeal, with no assurance that new difficulties will not emerge or that old ones will not find new expressions. No philosophical achievement – no argument, however decisive, no theory, however general – can solve all difficulties at once, prevent the apparition of new difficulties, or the resurgence of old ones. Philosophy is a constant state of remission, and so there is no sense in talking of a 'cure'.[46] And this has to do with Wittgenstein's conception of philosophy's subject-matter. That is, it is not an obstinate refusal of a cure that would be otherwise available, but the insistence that there cannot be any such cure. As I read it, it is precisely what Wittgenstein is suggesting in the *Zettel* when he says that "in philosophizing we may not *terminate* a disease of thought". It is true he then entertains the possibility of letting the disease "run its natural course" and of a "*slow* cure" (*Z* §382). However, the focus is not so much on the imagined state of health that would ensue, but on the fact that dealing effectively with philosophical 'diseases' is a long process – a process which might not be fully dependent on us. Ultimately, there is nothing like making a philosophical disease completely go away at once, no matter the seriousness of one's efforts. To be clear, I am not suggesting that the therapeutic dimension of quietism

[46] Once again, this can helpfully be contrasted with the early Wittgenstein's appreciation – or at least apparent appreciation – of what the *Tractatus* effectively achieves.

is exhausted by these considerations. Still, I would maintain that it gives us a way to appreciate its philosophical significance. If Wittgenstein was advocating for a definitive cure for all philosophical diseases – and so for a state of complete health – rather than for a series of therapies, then the contrast with traditional conceptions of philosophy might appear to be mostly verbal. After all, what would be distinctive about a therapeutic approach to philosophy then? It would seemingly yield the same kind of results as more orthodox approaches – or at least would similarly treat its conclusions (i.e. it would hope them to be not only all-encompassing, but also irrevocably established). The insistence on the variety of therapies, on the open-ended nature of our engagement with philosophical difficulties suggests an altogether different stance towards philosophy – its aims of course, but also the way it should be practiced.

It seems possible to understand the criticism of constructive philosophy along similar lines. Wittgenstein often opposes the descriptions that composes *Philosophical Investigations* to the explanations traditionally offered in philosophy (see for example *PI* §126). At the heart of this opposition lies a virulent criticism of our "craving for generality" (*BB* 17). Indeed, Wittgenstein appears to be of the opinion that orthodox approaches to philosophy, because they always aim for the greatest possible generality, end up reducing understanding – and especially philosophical understanding – to an act of subsumption. The particular, on these approaches, is something that needs to be explained away; it is only of interest inasmuch as it exhibits the general – which is, ultimately, the real object of investigation. And this is why philosophy traditionally conceives its task as one of revealing the essence of phenomena (e.g. following a rule, or language) – rather than appreciating the diversity of their manifestations (e.g. the heterogeneity of rule-following

activities or the many different connections that hold between language-games). Ultimately, Wittgenstein's point seems to be that, in so conceiving its task, philosophy fails to appreciate that another kind of understanding is within its reach (i.e. one that revolves around the 'seeing of connections'). Of course, attending to the many nuances among our rule-following activities, for example, to their various connections – in a word, to their family resemblance – is bound to appear dissatisfying to someone set on uncovering the essence of rule-following. But this only suggests that these are indeed different ways of conceiving philosophy's objects of investigations – and consequently, philosophy's tasks. On this point, Wittgenstein's distrust of our "craving for generality" very much echoes his criticism of the picture of philosophy as attempting to grasp infinite longitudinal strips. Indeed, a claim that would be absolutely general, would not accept any exception. And so, no new case could ever contradict it. It would be, not only all-encompassing, but also definitive.

Wittgenstein's discussion of rule-following in *Philosophical Investigations* – and in his later writings more generally – is rather instructive here. One striking feature is its constant engagement with a great variety of particular cases. There is the case of going on according to signposts (*PI* §§85, 87, 198), of reading off charts (*PI* §86), of grasping what a picture is a picture *of* (*PI* §§139-142), of continuing an arithmetic series (*PI* §§143-148, 151-155, 179-181, 185-190, 212-214, 226, 228) of reading aloud (*PI* §§156-171), of drawing a line in relation to another (*PI* §§174-175, 230), of grasping or understanding the meaning of a word (*PI* §§139, 191, 197), of playing or inventing a game (*PI* §§200, 204), of continuing a decorative pattern (*PI* §§208, 211), etc. Now, if Wittgenstein conceived rule-following as an infinite and indivisible longitudinal strip, and if he conceived his task

as that of grasping that strip (i.e. the essence of rule-following) *at once*, then the detailed examination of such a wide array of cases might appear superfluous – even by his own standards. But I am of the opinion that this is simply wrong. Wittgenstein really conceives rule-following as a series of cross-strips, as an open-ended series of variously related phenomena.[47] His attentiveness to particular cases is a constitutive – and inescapable – feature of his philosophical approach to rule-following.

I do not have the space to fully flesh out – nor to convincingly argue for – my proposed reading of Wittgenstein's discussion of rule-following, but we can at least say this: there is a temptation to read Wittgenstein's discussion as culminating in a number of maximally general claims about rule-following (e.g. "to follow a rule, to make a report, […] are customs" (*PI* §199), "'following a rule' is a practice" (*PI* §202), etc.). And there is a further temptation to take Wittgenstein's discussion as having the sole aim of establishing the truth of these claims, as if they were expressions of the *essence* of rule-following. Now, without denying the importance of these claims in Wittgenstein's discussion of rule-following – nor their apparent generality –, I am of the opinion that we should take seriously Wittgenstein's explicit focus on the ways in which cases of rule-following relate, but also differ from one another – and so potentially raise different concerns.[48]

---

[47] On this point, see, for example *RFM-VI* §21 where Wittgenstein discusses the circumstances in which we may say of someone that he "followed a rule (or a signpost; played a game, uttered a sentence, or understood one; *and so on*)" (my *emphasis*). The very way he frames this suggests that, for him, following a rule encompasses, not only a diversity of variously related phenomena, but an *open-ended* diversity.

[48] For example, it is quite clear that, for Wittgenstein, the notions of fit, ability and understanding are central to our discussion of rule-following. But he is also adamant that "the criteria which we accept for 'fitting', 'being able to', 'understanding', are much more complicated than might appear at first sight" (*PI* §182). And they are much more complicated, in part, because of the many ways in which something may *fit* something else (e.g. the picture of a cube and an actual cube versus the beginning of an arithmetic series and its continuation, etc.)

There is a danger to overstate the difference here. There is certainly a difference between, for example, proposing to uncover the essence of rule-following and carefully attending to the subtle nuances and connections between various rule-following activities. However, the transition from orthodox approaches to philosophy to quietism might be more a matter of degree. For one thing, it seems that, in addressing philosophical questions, one can claim to various degrees of generality. The choice is not between absolute generality and a complete refusal to generalise. There are many intermediary positions that appear to be perfectly cogent. If we come back to the picture of philosophy as occupied with either infinite longitudinal strips or their finite cross sections, we can frame the suggestion as follows: abandoning philosophy's traditional attempt to grasp the infinite longitudinal strips at once forces us to shift our focus to their finite cross sections. But it does not compel us to accept any one way to approach these cross sections. For example, if we can decide to only attend to their distinctiveness, we can also try to find patterns among them. That is, if we can be satisfied to only 'describe' these cross sections, we might also want to do more: to attempt generalisations, to formulate 'hypotheses' about them. Of course, adopting the latter approach probably reveals a certain 'craving for generality', but it might still be underpinned by a recognition of the tentative nature of its conclusions (generalisations or 'hypotheses') and ultimately of its open-ended nature. As such, it might be largely unproblematic – even in light of Wittgenstein's criticism of our craving for generality.

On this point, it is worth noting that even Wittgenstein's own aim of arriving at a surveyable representation partly involves being sensitive to similarities and patterns, in a word, to possible generalisations. The difference between 'seeing connections' among

cross-sections and forming hypotheses about how they relate to one another – between a surveyable representation and an explanation – might not be as fundamental as it is often assumed by the proponent of the former, if the forming of hypotheses is conceived as part of a continuous engagement with and careful attention to yet unconsidered cross-sections. The difference might not be clear-cut, but it does not follow that it is not genuine – or that it is not of philosophical interest.

What about McDowell's own commitment to quietism? As far as I can tell, this particular strand is largely absent from his early discussion of Wittgenstein. Not because, he is, at the time, purposely departing from Wittgenstein on this and advocating for his own alternative brand of quietism, but rather because he does not initially appreciate its importance in uniting the various expressions of quietism. As one of the rare advocates of quietism, this initial oversight has, I think, largely contributed to the confusion that surrounds it.[49] It also partly explains his greater tendency for generalisation (see especially McDowell 1984 and 1992). The difference between McDowell and Wittgenstein on this point also manifests itself in their attitude towards particular cases. Whereas for McDowell, they tend to serve as *illustrations* for the general points he is trying to make, for Wittgenstein, they are treated as the very objects of investigations. In contrast with Wittgenstein, McDowell tends to invoke very few cases and does so after having arrived at his conclusion.[50] We could say that, compared to Wittgenstein's, McDowell's actual practice of philosophy – at least in early papers –, is a lot more tolerant of a certain 'craving

[49] That being said, he seems to acknowledge this dimension of Wittgenstein's quietism in later writings (see for example McDowell, 1996, 177)
[50] See for example his discussion of mental content in McDowell 1992.

for generality'. But, as we have insisted, this greater tolerance does not necessarily amount to a relapse into an attempt to uncover essences.

In light of these considerations, we can tentatively propose the following conclusions: i) quietism is not an empty methodological commitment. It is neither unintelligible, nor does it collapse with constructive philosophy, although the difference might not always be as fundamental as it is often assumed by its proponents. ii) Wittgenstein is not obviously guilty of methodological dissonance, since making positive philosophical claims and laying down arguments is not, in itself, incompatible with the adoption of a quietist stance. McDowell's case is not as clear, in part because his view on the matter seems to have evolved. But he makes certain concessions in later writings that bring his philosophical stance more in line with Wittgenstein's. iii) Finally, there seems to be an important difference between two conceptions of constructive philosophy, one that sees its object as infinite longitudinal strips and another that sees its objects as a series of finite cross-strips. Whereas the distinction between the former and quietism is clear-cut; the distinction between the latter and quietism is far more subtle.

Our interest in quietism was twofold: we wanted reassurance that we could co-opt certain insights about rule-following from McDowell, without however committing ourselves to quietism. And crucially, that we could do so, without hereby distorting – or worse falsifying – these insights. In light of what we have said, I think there is no difficulty, in principle, if our attempt at constructive philosophy is of the unproblematic kind – and so, if we approach rule-following as a series of cross-strips, rather than as an indivisible and infinite longitudinal strip.

And this naturally leads us to our second concern: if we are not proposing to uncover once and for all the *essence* of rule-following, then *what are we proposing*? In what sense are we working towards a substantive account of the phenomenon, rather than a surveyable representation, in the way envisioned by Wittgenstein?

1.7 Philosophy Like Botany

Or to put the matter differently: what ultimately distinguishes a substantive account of rule-following from a quietist response to the regress of interpretations, in the form of a surveyable representation – given that an account also has to show us how the regress can be avoided and that a quietist response, just like a substantive account, says something positive about the phenomenon of rule-following?

In light of our previous discussion, it is clear that a substantive account of rule-following must aim at a certain degree of generality. After all, one would not be explaining what it is to follow a rule – and so would not be offering an account of the phenomenon – if one were to refuse to identify particular cases as instances of rule-following, to highlight common features of these cases and ultimately to make claims about rule-following in general. However, as we have seen, quietism – even in its most radical expressions – does not imply an inflexible hostility towards generalisation either.

As I see it, what matters most, is not so much the degree of generality reached, but rather the underlying *interest* in reaching any such degree. Ultimately, putting forward a substantive account of rule-following does reveal a greater *willingness* to generalise than the adoption of a quietist stance. Even more, it assumes a positive attitude towards attempts at generalisation – i.e. towards the formulation of hypotheses, the tentative grouping of

particular cases, on the ground of common features, etc –, in a way the adoption of a quietist stance simply does not. We could then say that what matters in deciding whether or not a series of positive claims about rule-following are an attempt at constructive philosophy, of the unproblematic kind, or a manifestation of quietism is whether or not, in making such claims, we give in to our *craving* for generality.

To work towards a substantive account of rule-following is to assume that the kind of intelligibility we reach when giving in to such a craving can be valuable, in itself – and so that the interest of our positive claims about rule-following is not solely determined by their success in escaping the regress of interpretations or in answering other philosophical difficulties.

There is a sense in which Wittgenstein's practice of philosophy in his later writings is always *reactive* – to perplexing pictures, to apparent paradoxes (like the regress of interpretations), etc. If he ends up making positive claims, it is always *in response* – and only in response – to such philosophical difficulties. Hence, for example, his insistence that philosophy "consists in marshalling recollections *for a particular purpose*" (*PI* §127, my *emphasis*). To adopt a quietist stance, is to not see the point in giving an answer, if no question has been asked, if there are not already some difficulties that call for an answer.

In a remark from his preliminary studies for the second part of *Philosophical Investigations*, Wittgenstein makes the point that classifications can be useful for "a particular practical purpose", but that he does not wish to offer such classifications, for the sake of it (see *LWPP-1* §727). He then adds: "A botanist classifies plants. But you don't need a system of classification to show somebody how multiform plants are and how diverse the fine distinctions among them are." (*LWPP-1*, §728). The analogy is, I think,

quite useful to illustrate our point: one can be in the business of classifying things without also being in the business of showing that they participate of a common *essence*. One can classify plants according to their genus and species, without assuming that one has uncovered, once and for all, the essence of what it is to be a plant. Still, classification – as it is done in botany or in other branches of natural history – remains an act of generalisation *par excellence* (or at least an act of subsumption *par excellence*).

Wittgenstein seems to suggest by this remark that his way of doing philosophy – and so his quietist stance – diverges from botany – or the *ethos* of botany – precisely in that it does not aim at a *system* of classification. As if, most of the time, it could contend itself with the mere recognition of the great diversity of phenomena encountered.

Keeping on with the analogy and in contrast with Wittgenstein's quietism, our suggestion then is that to engage in constructive philosophy, in an unproblematic way, is to adopt the *ethos* of botany. One can work towards a substantive account of rule-following by engaging with an open-ended diversity of particular cases of rule-following, just like one can work towards a system of classification of plants by looking at an open-ended diversity of species. In doing so, one is not hopelessly attempting to grasp at once the essence of the phenomenon. But one is not content with acknowledging the diversity of its manifestations either.

Chapter 2

Merleau-Ponty on Meaning and Motivation

2.1 Introduction

In the previous chapter, and in contrast with McDowell, we have committed ourselves to work towards an account of rule-following that is not only philosophically substantive, but also partially reductive. That is, we have proposed to show that it was possible to reconstruct meaning and understanding, as McDowell conceives of them, out of materials that do not presuppose such meaning and understanding. This would require, we have argued, to dig below the level of meaning and understanding McDowell takes to be bedrock, though not so low as to entirely deprive ourselves of intentional resources.

In this chapter we will suggest that a number of Merleau-Ponty's insights on perception and action can help us in that endeavour. We will focus especially on the conception of perceptual content and on the notion of motivation he lays out in the *Phenomenology of Perception*.

2.2 The Project of the *Phenomenology*

The P*henomenology* is first and foremost an investigation into perception. It is an attempt both to get clearer on the nature of perceptual content and on the ways that content can - and does - inform action and thought. Ultimately, it argues for a new *cogito*, for a new way of conceiving ourselves where we are shaped by and embedded in relations of

motivations. While a detailed commentary of the *Phenomenology* is not required for our current purpose, it seems, a brief exposition might be necessary to properly circumscribe the notion of motivation – and to appreciate both the role it plays in Merleau-Ponty's philosophy and the role it might come to play in our discussion of rule-following.

After all, it is not at all obvious that Merleau-Ponty has anything interesting to say on our current issue. As we have already recognised, the *Phenomenology* is not primarily about rule-following. And, on the face of it, it does not appear concerned – even in passing – with thematising the issue. With the notable exception of a brief section on the expressive powers of linguistic signs (see "Le corps comme expression et la parole"), the *Phenomenology* almost never explicitly discusses the kind of cases that are seen as paradigmatic of rule-following. One finds lengthy discussions of visual illusions, of odd psychosomatic phenomena, like phantom limbs or of the behavioural and cognitive deficits exhibited by brain-injured patients. But there are very few – if any – comments made about continuing an arithmetic series or a decorative pattern, being directed by a signpost, playing a game, etc.

And surely, this is something we will need to address in due time. If our proposal is to carry any weight, then the apparent incongruity between the queer phenomena discussed by Merleau-Ponty and the ordinary rule-following activities we are interested in will need to find some sort of resolution. Ultimately, our aim is to lay the groundwork of a satisfying account of rule-following – one that is partially reductive. But our proposal here is also partly exegetical: what we are claiming is that the notion of motivation *as* it is elaborated by Merleau-Ponty in the *Phenomenology* can play a central role in such a satisfying account of rule-following. We are advocating for some sort of *philosophical*

reuse. And so, evidently, it matters that what we end up transposing in the debate on rule-following is indeed faithful to what is sketched in the *Phenomenology*.[1]

The matter is further complicated by the fact that Merleau-Ponty belongs to a philosophical tradition that is largely ignored by the main authors who have shaped the debate on rule-following. And by the converse fact that Merleau-Ponty is not clearly aware of such a debate. In saying that, I do not want to perpetuate the popular prejudice that there is an insurmountable gap between Anglo-American – or analytic – philosophy and the so-called continental tradition. Quite the contrary. There might indeed be a gap; a persisting lack of dialogue between these two traditions – philosophy is not immune to insularity –, but our very proposal assumes that this gap, however imposing is not insurmountable. The point is more modest: Merleau-Ponty simply does not speak the same philosophical 'language' as Wittgenstein, his commentators and the authors who have later contributed to the debate on rule-following. And the difference is not just attributable to Merleau-Ponty's idiosyncrasies. The issues raised in the *Phenomenology* and, crucially, the way they are raised, are informed by the philosophical context in which Merleau-Ponty finds himself at the time. As they are, ultimately, by his choice of interlocutors.[2] He is attached

[1] The issue is not as straightforward as one might hope though. As we will come to appreciate in the following sections, even though the notion of motivation is central to the project of the *Phenomenology* – by Merleau-Ponty's own admission (see Merleau-Ponty, 1945, 76) –, its contour remains somewhat blurry. As we will argue, there are also persisting unclarities within the *Phenomenology* about how best to flesh out the notion – and with them, potential challenges. (see especially infra section 2.7). That being said, I still think, there's a meaningful difference between attempting to get clearer on Merleau-Ponty's philosophical insights and simply taking some passages of the *Phenomenology* out of context to fit our purpose. To get to the degree of precision and conceptual clarity that we seek, we might have to go beyond Merleau-Ponty. But to do so is not necessary to adopt a tendentious stance.

[2] We could say that Merleau-Ponty was in dialogue with Husserl, Heidegger, and Sartre – to name a few – in the same way that Wittgenstein was in dialogue with Frege, Russell, or Moore. His indebtedness to Husserl is evident from the very first pages of the *Phenomenology*. As for his engagement with Heidegger's and Sartre's respective works, it becomes more apparent in some of the later chapters dedicated to the question of time and that of freedom (see the aptly named "La temporalité" and "La liberté" in the third and last section of the *Phenomenology*).

to certain terms that have little resonance to analytic ears, while using other terms and expressions in ways that simply do not match analytic uses. Overlaps in philosophical concerns can thus be easily missed if we do not purposely attempt to find a common ground of expression. For example, Merleau-Ponty repeatedly suggests that objects of perception raise problems of "existential modality" (*modalité* existentielle) or of "transcendence" (see most notably Merleau-Ponty, 1945, 422-423). At first sight, the philosophical issue that Merleau-Ponty alludes to might appear altogether alien to our approach to rule-following – and ultimately to the broader debate. After all, nothing we have said so far has been framed in *these* terms. But, as we will argue, these passages in the *Phenomenology* can ultimately be read as a call for the kind of objectivity we have tried to circumscribe.

In other passages, Merleau-Ponty speaks of objects of perception as having a meaning. He suggests, for example, that the presence of objects in the visual field, what they are, or the way they are arranged can all have meaning *in perception*. Hence, interposed objects between us and the object to which we are attentive might have the meaning of distance. A rock on a walking path might have the meaning of an obstacle, etc. (see for example Merleau-Ponty, 1945). This is certainly at odds with the more circumscribed use of the term 'meaning' and its cognates in the analytic tradition. And so, we will need to tread carefully when this comes up.

How then are we to approach the *Phenomenology*, given our current concern? In the previous section, we have argued, with McDowell, that our responsiveness to meanings – or our understanding of signs – do not need to involve anything more than a perceptual or quasi-perceptual state. But then, in contrast with McDowell, we have suggested that there might be a way to account for such a responsiveness without invoking our being

perceptually sensitive to meanings themselves, but merely to some more fundamental – and yet to be specified – features of signs. A perceptual – or quasi-perceptual – state would thus still play a mediating role between signs and our response to them. However, it would allegedly not have the same kind of content as under McDowell's proposal, nor would it be internally related in the same way to signs and actions.

In light of this, I think our best way into the *Phenomenology* is through Merleau-Ponty's discussion of perceptual content – and the particular dialectic that guides it. This will seamlessly lead us to his discussion of the relation between perception and action – and enable us to then explore a number of his conceptual commitments when it comes to motivations.

2.3 The Content of Perception

What is the nature of the content of our perception? Merleau-Ponty does not frame the question in those terms, but I do not think doing so distorts the problem he is trying to circumscribe – nor his way of approaching it.[3] Indeed, when it comes to perception, Merleau-Ponty is interested in the mode by which things become available to us. His problem is not that of traditional epistemology – or at least it's not just that. His aim, in the initial sections of the *Phenomenology*, is to clarify the way in which things are *given to us* in perception – not to secure assurances that there is indeed correspondence between how things are given to us and how they really are. A lot more could be said here obviously, but the main point is this: for Merleau-Ponty, the question of what is given to us in perception is a self-standing philosophical concern. Obviously, the kind of answer we end up giving

[3] On this point, I am in complete agreement with Kelly (see Kelly, 2007, 23-25).

to that question will have various implications for the problem of traditional epistemology. However, our answer should not be influenced by any assumed constraints as to how the problem of traditional epistemology should be answered – nor as to how the relationship between mind and world should be conceived. So, instead of laying down from the start what should be given to us in perception *if* perception is to secure knowledge of how things are, we should, according to Merleau-Ponty, simply be attentive to what is indeed given to us in perception, without any undue prejudice.[4] Merleau-Ponty is keen to get clearer on what it is like, for example, to have the perceptual experience of an ashtray or of a red rug (see Merleau-Ponty, 1945, 433 and 499 respectively). And he ultimately thinks that this question can find an answer prior to and independently of the problem of the relation between how the ashtray or the rug manifest themselves in perception and how they really are, as a matter of fact.

So then, how should we conceive the content of our perception? An initially plausible suggestion might be that it corresponds, more or less, to what is often called 'pure sensations' or 'sense-impressions' (see Merleau-Ponty, 1945, 26). What is given to us in perception are these sensations of round and red, of shapes and colours *before* we conceive of them as sensations of something being round or something being red, etc. Upon gently pressing my closed eyes, for example, I have a grey visual impression that slowly turns into a red one. And this later visual impression is not, in itself, an impression *of* anything. Nor is its enjoyment – its being given to me – dependent upon me subsuming it under the concept *red*. The suggestion then is that the same goes when I perceive the ashtray or the

[4] On this, the overall structure of the *Phenomenology* is rather telling: the new *cogito* concludes the *Phenomenology*, rather than being a necessary foundational block to the project. It is a complete inversion from the role the *cogito* plays, let's say, in Descartes' *Meditations*.

red rug: the content of my perception can ultimately be broken down into a series of such impressions.[5]

Merleau-Ponty however takes issues with such a picture of perceptual content. To begin with, he insists that, as a matter of fact, we never encounter in perception anything that would correspond to sense-impressions – or pure sensations. After all, implicit to the notion of sense-impressions is that their enjoyment is ultimately independent of the enjoyment of any other impressions. For the red visual impression to be given to me upon pressing my closed eyes, I do not need to enjoy or have enjoyed any other perceptual content. In the case of the red rug, my perception of it might involve a number of distinct sense-impressions, but my enjoyment of any one of these sense-impressions taken in isolation would not be dependent upon anything outside of itself. However, as Merleau-Ponty insists, when I perceive the red rug, the red appears inseparable from the texture of the fabric. As he puts it: in the case of the rug, there is something woolly about the red, without which it wouldn't be, to us, the *same* red (see Merleau-Ponty, 1945, 27). Even for the red I perceive upon pressing my closed eyes: it seems it has the quality it has because of its slightly brighter contour, of its peculiar 'vibration'.

And so, for Merleau-Ponty, this seems to raise the question: what can ever ground one's appeal to pure sensations - or sense impressions - to account for perceptual content? Why should we posit such entities, in principle, if we never enjoy them in actual perceptions? Ultimately, Merleau-Ponty is adamant that we should be wary of such a move. For him: "only the effective structure of perception can teach us what it is to perceive"

[5] I think it would be fair to assimilate such a view on perceptual content to the kind of views McDowell – and before him Sellars – criticise as falling prey to the myth of the given (see McDowell, 1996, especially lectures I and II and Sellars, 1997).

(Merleau-Ponty, 1945, 26).[6] To try to uncover the content of perception underneath actual perceptions is simply misguided.[7]

But Merleau-Ponty goes further. For him, the most fundamental flaw of such a picture of perceptual content is that it has difficulties accommodating the 'meaning' which appears to adhere to or be constitutive of our actual perceptions (see Merleau-Ponty, 1945, 26 and 36). The claim is repeatedly made by Merleau-Ponty, but its implications are far from clear at the outset. One potential confusion comes from Merleau-Ponty's reliance on the notion of structure or form (*Gestalt*), which he borrows from psychologists like Koehler and Kofka. Indeed, Merleau-Ponty is keen that our actual perceptions are always organised according to structures, like that of figure-ground. And so, when we perceive a thing (e.g. a red rug) we always perceive it as a figure that detaches itself from its immediate surrounding (e.g. the hardwood floor on which it is laid) and so presents itself to us as on a background. Following Koehler and Kofka, Merleau-Ponty insists that such structures or forms are irreducible to the sum of their part. That is to say, for him – as for them – the

[6] My translation. All further citations of Merleau-Ponty, except otherwise indicated, are also my translation of the French original.

[7] This is obviously not a decisive objection. Merleau-Ponty seems to concede that much, and to recognise that an advocate of the pure sensation picture – or as Merleau-Ponty calls it "the empiricist picture" – of perceptual content might simply refuse to make such a concession. In which case, Merleau-Ponty's appeal to the woolly red or of the peculiar red of closed eyes might be mute. For the advocate of the pure sensation picture, they would just be perceptual phenomena that need to be further analysed and broken down. The same would go with Merleau-Ponty's appeal to certain results in psychology. For example, Merleau-Ponty takes it to be well established that the most basic contents humans and other animals are *responsive to* in perception are not mutually independent sense-impressions, but structures or forms (*Gestalt*). But then again, it is not obvious that an advocate of the pure sensation picture could not explain away these results… The point is that it is only from a particular approach to perception that the pure sensation picture of perceptual content becomes plainly at odds with scientific evidence and with what we might call phenomenological evidence. Merleau-Ponty thinks that such an approach will ultimately prove more fruitful, but, in itself, this does not amount to a refutation of the pure sensation picture (see especially Merleau-Ponty, 1945, 46, 47 and 49. It does not matter greatly for our current concern, but I do think this might constitute a point of friction with his approach in *The Structure of Behaviour* (see Merleau-Ponty, 1942).) There is a further question: independently of Merleau-Ponty's way of framing the issue, might an advocate of the pure sensation picture simply – and coherently – recognise structures or forms (*Gestalt*) as possible pure sensations? And so, avoid the need to 'explain them away'? Committed sense-data theorists, like Ayer, might offer an interesting example of such attempts (see for example Ayer, 1945, 304-305).

actual perception of the red rug on the hardwood floor cannot be analysed into the perception of bits of woolly red on bits of polished wood. However, I think we should resist the temptation to understand Merleau-Ponty's talk of perceptual 'meaning' as simply amounting to an appeal to that particular notion of structure or form. Indeed, I would contend that, in speaking of perceptual 'meaning', Merleau-Ponty does not merely claim that there is something that is given to us in perception, which is irreducible to the sum of its part. His suggestion here is really twofold. First, he argues that *what is given to us in perception* are structures or forms. Second, he suggests that these structures are, in some sense, meaningful. And it is not at all obvious that accepting structures as the basic content of perceptions forces us to accept them as being meaningful.[8] It is not to say, of course, that there might not be reasons as to why structures – in contrast let's say with pure sensations – can be and actually are meaningful. The point is simply that the notion of meaningfulness does not seem to be implied by that of structure.[9]

Moreover, as we have already hinted at, another potential confusion comes from a mismatch between the notion of meaning Merleau-Ponty appeals to here and the narrower conception we find in the analytic tradition. Indeed, Merleau-Ponty claims that the

---

[8] As a matter of fact, Merleau-Ponty's stance towards Koehler and Kofka is precisely that they have recognised the importance of structures or forms in perception, but then have failed to recognise their meaningfulness – or at least the implications of their meaningfulness (see Merleau-Ponty, 1945, 73-77 and Merleau-Ponty, 1942, 207-219). Merleau-Ponty does consider that they have missed something important there, but I think it would be farfetched to say that their stance on the issue is ultimately incoherent.

[9] If this is right, then the notion of meaning appealed to here would also not be dependent upon any prior commitment about *the role* of structures or forms in perception. This especially matters, I think, when one tries to argue for the contemporary relevance of Merleau-Ponty on these issues. If his philosophical insights about the meaningfulness of actual perceptions were to rest, not only on empirical evidence, but on empirical evidence as they were theorised by psychologists like Koehler and Kofka, then it would matter a great deal whether or not these theorisations have survived the tribunal of experience. If, on the contrary, Merleau-Ponty's insights can be reconstructed independently of any such commitment, then we can largely ignore any dispute about the empirical standing of the so-called psychology of forms.
This, of course, does not diminish the importance of the notion of structure in the genesis of Merleau-Ponty's thought, nor its importance in the dialectic he sets forward in *The Structure of Behaviour* and at the beginning of the *Phenomenology*.).

perception of the red rug, inasmuch as it is given to us as a figure on a background, possesses a meaning. And while he does not explicitly tell us what that meaning amounts to – at least not at first –, it is clear that our perception of the red rug does not stand for anything in the way the word "red" might stand for *red* let's say. Or again, coming back to the example of the rock on the path *meaning* an obstacle, it is clear that the relation there is not that of a conventional sign and its agreed upon referent. So then, in what sense are these meaningful? A tempting move might be to say that these perceptions are meaningful inasmuch as their content is conceptually articulated.[10] So the perception of the red rug is meaningful because its content is not a series of pure sensations but is already informed – or structured – by the concepts *red* and *rug* – and ultimately by other very general concepts like that of *figure on a background*. In the case of the rock on the path, it might be that the content of such a perception involves the concept *obstacle* – and it's in that sense that the former *means* the latter.

Some might find such a line of thought appealing. However, Merleau-Ponty is certainly not one of them. Indeed, after dismissing any attempt to ground perceptual content on pure sensations, Merleau-Ponty briefly considers a suggestion along these lines, only to quickly dismiss it. The main issue for him is that our perceptions – and their content – are ultimately not as *determined* as our concepts.[11] Take the red rug again. If what was given to us in

[10] The move is certainly one that has been made in the analytic tradition. McDowell's project in *Mind and World* is a clear example of it (see McDowell, 1996). However, as we will see, such a position can also be articulated from the other side of the analytic/continental divide.

[11] It is worth noting that Merleau-Ponty raises a very similar issue with conceiving perceptual content in terms of sensory qualities rather than in terms of sense impressions. Indeed, he briefly considers the suggestion that the red of the rug might be given to us in perception not as a mere impression of red, but as the quality the rug has of being red. In which case, the content of our perception would somewhat correspond to its object. What would be given to us when we perceive the red rug would be the rug itself as having the qualities it has, as a matter of fact (e.g. the quality of being red, woolly, etc.). Now, Merleau-Ponty does not raise this point, but, as I see it, such a suggestion appears ultimately committed to some form of naïve realism – or relationalism. However, in contrast with McDowell's own brand of naïve realism, what would be given

perception was *the rug being red*, then our perception would have the same content as our judgement that the rug is red. And if the concepts *red* and *rug*, etc. are involved in such a judgement, then they would also be involved in the corresponding perception *in the same way*. The issue, for Merleau-Ponty, is that actual perceptions have a certain indeterminacy that judgements simply do not have.[12] The boundaries of our concepts – as they are involved in the latter – are sharper, more precise than the meanings we find in perception. The point here is not that the content of our perception is somewhat more fine-grained than the concepts we employ in judgement – and so that they somewhat always fall through.[13] The point rather is that what is given to us in perception is not something on which we have *a full grasp*. Merleau-Ponty illustrates this by appealing to our difficulty in perceiving a

---

to us in perception here would not be already shaped by concepts (see McDowell, 1982 and especially 2013d.). The suggestion might be more akin to the kind of naïve realism defended by Brewer or Campbell (see Brewer, 2011 and Campbell 2002). In any case, Merleau-Ponty dismisses the suggestion on the ground that what is given to us in perception is simply not as *determinate* as the qualities of actual objects. Merleau-Ponty sees visual illusions as perfect illustrations of this point. Take for example the Müller-Lyer visual illusion, where a line appears longer than another one, despite being actually of the same length. Clearly, as actual objects, both lines are of a determinate length, but if they were really given to us in perception in that way, why would we so systematically 'mis-see' them. For Merleau-Ponty, the answer is obvious: in the objective world, lines are of determinate length. But that's not so in perception. In perception, both lines become integrated in a meaningful structure or form and that's how they are ultimately given to us (see Merleau-Ponty, 1945, 28). This clearly mirrors his stance on the suggestion that the content of our perception is modelled on that of our judgements.

[12] Another related issue for Merleau-Ponty is that in conceiving what is given to us as already conceptualised in that way, we are in danger of eroding the distinction between perception and judgement (see Merleau-Ponty, 1945, 58-59). Indeed, for him, there is a "prise de position" in judgement which is absent perception. There is, we could say, a fundamental *passivity* to latter, for which we must be able to make room (hence phrases like "being *given* in perception" or "appearing in this or that way"). As, I see it, such a line of reasoning might be resisted – and actually has been resisted most notably by McDowell who maintains that we can recognise both that perceptual content is always already conceptually structured and that the delivery of such content can still reveal a certain *passivity* (see McDowell, 1996). Since this issue is not crucial to appreciate Merleau-Ponty's position nor in need to be settled in light of our current concern, I will simply avoid taking a stance on it.

[13] McDowell, for example, attributes such an argument to both Evans and Peacocke. It is an argument that he ultimately deems inconclusive (see McDowell, 1996, 56-60). But this is not the argument Merleau-Ponty puts forward. (On this point, Kelly's discussion of the dispute between McDowell and Peacocke is enlightening. *Contra* McDowell, Kelly is sympathetic to the thought that perceptual contents are non-conceptual, but he does not see the alleged fine-grainedness of such contents to be the most relevant argument in favour of it. For him, what matters most is what he calls the "situation dependence" of perceptual contents, which he ultimately associates with Merleau-Ponty's view in the *Phenomenology* (see Kelly, 2001, See also Peacocke 1998 and McDowell 1998).

square *as being square* when it is resting on one of its corners rather than one of its side (see Merleau-Ponty, 1945, 73). If our perception was indeed shaped by the concept of square, then the orientation of the figure would not matter at all, not just in the judgement we ultimately make of its shape, but also in the actual perception we have of it. In the same way, the concepts of up and down are very clearly relative to orientation. However, while we have no difficulty in conceiving that the same square could rest on one of its sides as well as on the opposite, a scene seen upside down simply does not appear to us to be the same as the scenery seen upright. If geometrical figures easily lead themselves to formal manipulations like that of rotation in abstract reasoning, the same cannot be said of our actual perceptions. Of course, for Merleau-Ponty, this speaks in favour of the content of perceptions being structures or forms – in that the actual perception of the square or the scenery is dependent upon the broader perceptual context. But in bringing these cases up, Merleau-Ponty is also opposing the thought that what appears to us in perception could only ever be meaningful if it had the kind of exactness that underlies the subsumption under a concept and ultimately makes possible abstract reasoning.[14]

Now, if the meaning Merleau-Ponty assumes in perception could only be understood in that way (i.e. as dependent upon our conceptual – and ultimately rational – capacities), then we would be facing two unsavory alternatives: either to reduce perceptual content to pure sensations or sense impressions or, on the contrary to see perceptual content as always already informed – or shaped – by our conceptual capacities. For Merleau-Ponty, the error of the first proposal is to not give itself the resources to account for that meaning. The error of the second is to make that meaning too determinate. Ultimately, Merleau-Ponty sees this

[14] The point about abstract reasoning is a general one for Merleau-Ponty. It applies to the kind of geometrical manipulations we just discussed, but also to deductive reasoning.

apparent dilemma as emerging from a flawed dialectic between an empiricist picture of actual perception and an 'intellectualist' – or rationalist – one. It is flawed because it assumes perception to be fundamentally aimed at representation, at offering us a view of how things are, and in doing so misses the phenomenon of perception (see for example Merleau-Ponty, 1945, 27 and 50). For Merleau-Ponty, the way out of the dilemma is to make room for this still elusive meaning in perception

It is at this point in the *Phenomenology* that the notion of motivation first comes into play. The empiricist and intellectualist pictures, despite accounting for perceptual content in diametrically opposed ways, are ultimately two sides of the same coin. They both fall prey to what Merleau-Ponty variously calls the "prejudice of the world" or the prejudice of "objective thought" (Merleau-Ponty, 1945, 27 and 50). That is, in both assuming that perception, when successful, is essentially the unfolding of how things are to consciousness, in a disinterested act of knowledge, they prejudge the nature of the relation between world and mind.[15] For Merleau-Ponty, they ultimately commit themselves to a dichotomy: either they assume that the relation between world and mind is merely causal – or they maintain that the relation is irreducibly rational. In the first case, we might end up with some versions of the empiricist picture in which things in the world cause in us pure sensations which are then organised by various psychological laws of association. In the second case, we rather end up with some versions of the intellectualist picture: things in the world being given to us in such a way that they can serve as a reason, not only for our judgements about them, but for our very perceptions.[16]

[15] This echoes the point made earlier about the relation between the question of the nature of what is given to us in perception and what we have called the epistemological question.
[16] The psychology of form would be, in that respect, a limit case. Indeed, while denying that pure sensations constitute the content of our perceptions, it still maintains that the relation between how things are and how

The way out of the dichotomy, for Merleau-Ponty, is to abandon this prejudice of objective thought and to take seriously the suggestion that the world reveals itself in perception as already meaningful, not to a disinterested observer – Merleau-Ponty might talk of an ahistorical cartesian ego –, but to an embodied and situated agent (see Merleau-Ponty, 1945, 79). The notion of motivation, for Merleau-Ponty, can ultimately occupy the conceptual space thus opened and so offer us an alternative between reasons and mere causes. The suggestion, in a nutshell, is that it is possible to conceive of an action as being not only informed but also guided by perception, in a way that goes beyond the relation that holds between an effect and its cause(s), without however depending on the involvement of any properly rational capacities (as they manifest themselves in concept application and reasoning).

Thus, we can now better understand Merleau-Ponty's insistence that the question of the nature of perceptual content is self-standing *with regards to the problem of traditional* e*pistemology*. The very conceptual framework in which this problem finds its expression is precisely what Merleau-Ponty thinks falsifies the phenomenon of perception. But while the question is self-standing in that respect it now emerges that a satisfying answer cannot, for Merleau-Ponty, abstract from the role of perception in guiding action. We were unable to clearly circumscribe the notion of meaning Merleau-Ponty claims is indigenous to perception. And this is partly why. From what we have said so far, it might have become clear that, for Merleau-Ponty, the content of perception is meaningful in that, it is intrinsically *intentional*. But the suggestion remains somewhat obscure. The way to make

---

we perceive them *must* be causal (see Merleau-Ponty, 1945, 75-76). This shows, once again that, while inherited from the psychology of form, the notion of structure or form plays an original role in Merleau-Ponty's philosophy.

progress is now given to us: the meaning – or intentionality – indigenous to perception is indissociable from action.[17] And so, if it ultimately grounds relations of motivations, it is also in them that its features are bound to be made most salient.

What remains to be seen is whether the notion of motivation, as it is thematised by Merleau-Ponty is indeed a genuine alternative to the apparent dichotomy between reason and mere cause.

2.4 Perception and Action

So, for Merleau-Ponty, we are missing something essential about perception – and especially about the content of perception – when we conceive of its relations to action in purely causal terms. But we are not getting things right either when we conceive the relation between perception and action as always mediated – or structured – by our properly rational – or conceptual – capacities. In the *Phenomenology*, Merleau-Ponty discusses in detail two empirical cases, which, he thinks, reveal the inadequacy of such ways of conceiving the relation between perception and action: the case of phantom limbs and that of agnosia.

Whether or not Merleau-Ponty is right on the details of these empirical cases does not matter greatly for our current concerns. And so, if we discuss them now, it is not so much as evidence in favour of Merleau-Ponty's conception of the relation between perception and action, but rather as illustrations of it.

When it comes to phantom limbs, Merleau-Ponty takes interest in a number of recurring clinical manifestations in amputated patients (e.g. persisting sensations 'in' the

[17] Hence Merleau-Ponty's remark on action – and more generally on motricity – that "what we have discovered by the study of motricity is, ultimately, a new meaning for the word 'meaning'" (Merleau-Ponty, 1945, 182).

amputated limb, obstinate attempts to realise tasks with the use of the lost limb, refusal to explicitly acknowledge the loss of the limb, etc.). The question is: how can we make sense of such clinical manifestations? One way would be as follows: the patient is ultimately conscious of the fact that they have been amputated, but, that consciousness is painful and so they are repressing it. In this process of repression, they end up forging the false belief that the amputated limb is still there. Their obstinate attempts at using the lost limb and their refusal to acknowledge its loss are simply behavioural manifestations of that false belief. This would be a "psychological explanation" of phantom limbs. For Merleau-Ponty, the issue with such an explanation is that it cannot account for the fact that the surgical section of the limb's afferent nerves often leads to the disappearance of the clinical manifestations of the phantom limb (see Merleau-Ponty, 1945, 105). After all, how could a physiological change in the peripheral nervous system interfere with the psychological process of repression – and even more, undermine the false belief in the presence of the limb?

The temptation then, may be to offer a "physiological explanation" of phantom limbs. The clinical manifestations would not be underpinned by a psychological process of repression, but rather by a purely physiological process, like the misfiring of the afferent nerves of the amputated limb. However, for Merleau-Ponty, this will not do either: such an explanation implies that conscious thoughts are not really involved in bringing about clinical manifestations of phantom limbs. However, this seems to be contradicted by the fact that memories and other conscious states often trigger sensations 'in' the amputated limbs, and especially sensations of pain (see Merleau-Ponty, 1945, 105).

The issue, I think, can be summarised as follows: psychological explanations reduce phantom limbs to a series of thoughts – or contents of consciousness – and relate them as reasons. The realisation that one has lost one's limb justifies – or serve as a reason for – the psychological distress. In turn, the distress serves as a reason for the false belief in the presence of the limb, which then serves as a reason for the attempt to use the lost limb or for the explicit denial of its loss. Physiological explanations, in contrast, reduce phantom limbs to merely causal relations between physiological structures (e.g. from the affected afferent nerves, to the higher order neurophysiological structures of the central nervous system, to the efferent nerves, whose activation bring about the behavioural manifestations). The issue then, for Merleau-Ponty is that, by each standing on their respective side of the reason/mere cause divide, these two kinds of explanation ultimately miss the phenomenon of phantom limbs.[18]

Instead, Merleau-Ponty suggests that, in the case of phantom limbs, the various content of consciousness and behavioural manifestations that are characteristic of the phenomenon are rather related by *relations of motivations*. The sensations 'in' the lost limb are not a mere collection of sense impressions, but they do not have the determinate content of a judgement either. They reveal an awareness of a particular way of being in the world – Merleau-Ponty talks of "existential attitude". That is, an enduring awareness of the

[18] As I read him, Merleau-Ponty does recognise that a thought – or any other content of consciousness – could serve as a reason for another thought *and* somewhat bring about or cause that other thought. As he might put it: a content of consciousness might *necessitate* another such content, like a "*cogitatio* necessitates another *cogitatio*". But we might, in the same breath, also talk of a "causality of idea" being at play (Merleau-Ponty, 1945, 115) and so of the former content causing the latter. For Merleau-Ponty, the important point, as I take it, is that while physiological and psychological explanations might be, on their own, inadequate, attempts to appeal to both mere causes and reasons are simply unintelligible. Now, obviously Merleau-Ponty tends to oppose "causes" to reasons – rather than "mere causes". But if we are right in our reading here, the contrast is really between mere causes and reasons, which may also be causes.

behavioural responses that were once made possible by the limb – and relatedly the practical meanings, in one's surroundings, that used to elicit such responses. For Merleau-Ponty, it is clear that such an awareness is in part dependent upon the workings of the afferent nerves, that keeps "live" such practical meanings and patterns of responses. But it is not reducible to the working of these nerves (see Merleau-Ponty, 1945, 116). After all, for Merleau-Ponty, the sensations 'in' the lost limb further crystallises meanings of loss – in the same way that memories of the accident or negative emotions towards the loss also give expressions to such meanings. And surely these meanings cannot be reduced to the working of afferent nerves!

Merleau-Ponty's suggestion is that the persisting sensations 'in' the lost limb, the memory of that loss or the emotion towards it can motivate one another. It is not merely that one content of consciousness can bring about the others (or vice-versa), but that each of them are somewhat internally related to the others, on the ground of their meanings (see Merleau-Ponty, 1945, 115-116). The same goes for the behavioural manifestations of the phantom limb: if the sensations 'in' the lost limb brings about obstinate attempts to use it or even the denial of its loss, it is because of their meanings. The relation between the sensations and the behavioural manifestations is like that between the first few notes of a melody and its continuation. The behavioural manifestations do not merely follow the sensations; they *fit* them (see Merleau-Ponty, 1945, 107).

Merleau-Ponty arrives at similar conclusions when he discusses cases of agnosia – and the apparent behavioural and cognitive deficits that comes with it. He is especially interested in case studies where patients reveal themselves competent in engaging practically with familiar objects and situations, but consistently fail to engage with the

same objects and situations in a detached way. A patient for example, reveals himself unable to identify the shape of an object put in his hands, or the relative position of his hands in relation to the rest of his body. He also reveals himself unable to perform simple 'abstract' movements upon request (i.e. to raise his arms, or to flex his fingers). Still, he is able to perform the same movements when a practical task he is familiar with requires it (see Merleau-Ponty, 1945, 132-133). Thus, he cannot identify the shape of an object, nor the position of his hands and he cannot flex his fingers in abstraction from any object or situation, but, in his workshop, he is able to flex his fingers to grab a hammer or any other tools he is familiar with. He is even able to perform rather complicated work, with almost as much dexterity as before the onset of agnosia (see Merleau-Ponty, 1945, 135-136).

According to Merleau-Ponty, such case studies are rather perplexing if we think they need to be explained either in terms of reasons for actions or in terms of mere causes (see Merleau-Ponty, 1945, 154-156). It seems such a patient is unable to form judgements about the shape of objects around him or about the position of his own body.[19] And so, the work he performs in his workshop cannot be explained in terms of such judgements guiding his hands in the performance of the work. But at the same time, the ease with which the patient performs the work cannot be imputed to a series of acquired reflexes. More is at play than the blind working of physiological mechanisms: the patient is ultimately sensitive to the practical meaning of his tools. The perception of the hammer does not merely cause his performance of the work, it also motivates it.

Now, for Merleau-Ponty, such considerations apply more generally. For him, it is not just pathological cases that resist explanation in terms of reasons or in terms of mere causes:

[19] The patient certainly cannot give verbal expression to such judgements, but Merleau-Ponty is of the opinion that it is not merely a verbal deficit, but a more general deficit of representation.

it is also our normal practical dealings with the world. When we practically engage with our surroundings, the relation between our perception of these surroundings and the actions we perform in them is also best conceived in terms of motivation. Cases of phantom limbs and of agnosia simply illustrate quite vividly the need for – or at least the fruitfulness of – such a shift in conception.

Obviously, Merleau-Ponty's discussion of phantom limbs and of agnosia is far richer than this brief summary might suggest.[20] But in the current context, it should suffice. It lays down the necessary conceptual apparatus to discuss two potential issues with Merleau-Ponty's general view on the relation between perception and action – issues which are deeply relevant to our underlying project –, namely the richness of perceptual meanings and the risk of a layered conception of our engagement with the world.[21]

2.5 The Richness of Perceptual Meanings

We have insisted that, for Merleau-Ponty, perceptual forms are meaningful and can be given to us in such a way as to guide action, without the involvement of properly rational or conceptual capacities. There might be a temptation to see this as imposing a constraint on the – intentional – richness of perceptual forms. That is, we might assume that the meanings indigenous to perception, inasmuch as they logically and developmentally precede our full-fledged mastery of concepts, ought to be quite simple and closely aligned to our basic behavioural capabilities. After all, according to a rather uncontentious

[20] For one thing, it suggests a number of other ways in which it may be important to distinguish normal and pathological dealings with the world (see especially Merleau-Ponty, 1945, 160-174).

[21] This way of framing things might give the impression that these issues are wholly independent of one another. I do not think this is the case, quite the contrary. However, exploring the many different ways in which these issues cross paths in Merleau-Ponty's writings would take us too far off our current concern.

philosophical view, if I can entertain the thought that this is a hammer and so that it is a tool designed to be used in this or that way (i.e. if I can entertain a thought with the content *this is a hammer…*), it is because I master the relevant concept *hammer*.[22] And so, it is difficult to see how a perception of the hammer could have *that* content if no concept whatsoever were involved in shaping it.

But what about my perception of the hammer as being *graspable* in *this* or *that* way? Can my perception of the hammer have that content, be meaningful in that way, without assuming that I have conceptually structured representations of the hammer and of my own body – and so ultimately without assuming that I master the concepts that would be involved in such representations? And what about my perception of the hammer as fitting a particular kind of work or task? Now I think many will be happy to concede that my perception of the hammer could be the perception of a graspable something – without any concepts being involved either in structuring the content of my perception or in informing the various steps leading from that perception to my effectively grasping the hammer.[23] (Although, many will see such a concession as implying that the relation between my perception and action is then only merely causal and so that my perception is not *guiding* my action in any interesting sense.) The matter will be a lot more contentious I think when it comes to the question of whether or not I can perceive the hammer as an object with a particular purpose – and so as being internally related to a particular kind of work or task. It will depend partly on what kind of work we consider reasonable to assign

[22] I assume here that the mastery of the concept of hammer implies more than just the ability to appropriately discriminate between hammers and non-hammers. As I see it, it also implies the ability to situate it within our broader conceptual space. One's mastery of the concept of hammer is dependent upon one's mastery of related concepts of tool, of work, etc.

[23] This is essentially what Merleau-Ponty suggests in discussing the case of patients with visual agnosia (see especially Merleau-Ponty, 1945, 134).

to the hammer. Surely, there is a clear sense in which the hammer can be used to hit objects with a particular motion of the arm. But can I perceive the hammer as fitting such a motion – and ultimately such an action - without already conceiving the hammer *as a hammer* and the work at hand as one of *hitting*? And what about more complex works? Isn't there a sense – maybe particularly salient to woodworkers – in which the hammer fits the nailing of pieces of woods together? Or for shoemakers a sense in which the hammer in the workshop fits the shaping of a piece of leather on a wooden last? Surely in both cases such *judgements* about the hammer could be made. But can the hammer be *perceived* as being so fitted without supposing antecedent – and conceptually articulated – beliefs about hammers and their specific involvement in woodworking or shoemaking?[24]

There is an initial uneasiness, I think, with the suggestion that the hammer could be perceived as a tool fitting a particular work (i.e. could have that perceptual meaning), without the involvement of our conceptual or rational capacities. And as the complexity of the work grows, so does the uneasiness. Of course, the most obvious way to escape the uneasiness here is simply to deny that, under such conditions, the hammer could have that perceptual meaning. (As I see it, this is very much what one does when one insists that the relation between perception and action ought to involve some form of practical reasoning – and so that the only way to get, for example, from the shoemaker's perceiving the hammer to his using it to make shoes is by invoking a series of intermediary steps involving conceptually articulated beliefs about the hammer, the leather and the last as well as

[24] In what follows, I will talk of the hammer as *fitting* a particular work, but it is worth noting that Merleau-Ponty often rather talks of tools as *calling for* a particular work (see Merleau-Ponty, 1945, 136). I will avoid this and similar expressions because, I think they raise – or at least make especially salient – a further issue, which we will only address in section 4.3.3, namely the extent to which meanings indigenous to perception can have *normative force*.

intentions on the part of the shoemaker.) But I think, such an escape is only required if it can be shown that our uneasiness is warranted. It would be bad practice to simply assume that it is.

So what is the source of our uneasiness at the suggestion that, in the workshop, the hammer is so given in perception to the woodworker or shoemaker? The source, I think, is twofold – and registers two distinct, albeit interrelated, worries. First, there is the impression that such a content is simply too intentionally rich to not be structured – or shaped – by rational – or conceptual – capacities. Second, there is the impression that such a content is simply too distant from their basic engagements with the world for it to be obvious that it can be guiding their respective action.

(i) Intentional richness and Conceptual Mastery

Now, what might give us the impression that such perceptual meanings are too intentionally rich to not be conceptually articulated? Might it be the fact that, in discussing the hammer as given in perception as a tool fitting a particular kind of work – that is, in discussing the hammer as having *that* perceptual meaning – I am ultimately appealing to and relying on a rather complex conceptual scaffolding? But surely, this is merely a testament of what it is to describe or explain something and so has very little to do with the phenomenon actually described or explained. One might have a very complicated story about what is involved in the perception of food sources by chicken or in the perception of potential mates by hoverflies.[25] But it would be a mistake to conclude that the conceptual

[25] The first case is discussed by Merleau-Ponty in the *Structure of Behaviour* and alluded to in the *Phenomenology* (see Merleau-Ponty 1942 and 1945). The second case, of course, is famously discussed by Millikan (see Millikan, 1990).

refinement required to tell these stories ought to actually structure the perceptual capacities of the chicken or the hoverflies. So of course, in *saying* that, in perception, the hammer can have the meaning of a tool fitting a particular kind of work, I am appealing to the concepts of *hammer*, of *tool*, of *work*, etc. And in *describing*, for example, the shoemaker as shaping a piece of leather on a wooden last, I am also appealing to concepts pertaining to shoemaking (e.g. of last). But this is not to say that to *perceive* the hammer in such a way – and to *use* it as such in shoemaking – requires the involvement of these concepts.

Maybe our uneasiness towards recognising that meanings given in perception can be intentionally rich even if they are not conceptually articulated rather comes from the impression that such a move would commit us to the unpalatable view that there could then be perceptual meanings echoing complex practices and yet completely independent of the concepts that inform – and ultimately make possible – these very practices. Take again our shoemaker. Couldn't one argue that we could not make shoes – at least not in the way the shoemaker makes them – were it not for a number of conceptual achievements? For example, couldn't one argue that the very work of shaping a piece of leather on a wooden last would simply be out of reach of people, altogether lacking *some* conception of what it is to make shoes – and so ultimately lacking *expressible* (we might alternatively say *shareable*) knowledge of the various steps involved in realising such a work, of the tools required, etc.? Shaping a piece of leather on a wooden last is not the same thing as randomly hitting a piece of leather on a piece of wood.[26] It is an action – or work – that can only be performed by people who somewhat understand what it is to make shoes (including the

[26] Of course, if one were to deny that there is indeed a *real* difference between shaping a piece of leather on a shoe last and merely hitting it, the suggestion that my perception of the hammer could fit the former – in a way that would go beyond merely fitting the latter – wouldn't even get off the ground.

role of hammer, the leather and the last in the process, etc.).[27] And how could reaching such an understanding not be dependent upon the mastery of concepts? [28]

One might take these considerations to clearly speak against the suggestion that meanings relating to shoemaking could somewhat be given in perception, without the involvement of properly conceptual capacities. But I do not think this is quite right. It is not at all clear that we should accept that there is indeed such a dependency of practices (like shoemaking) on conceptual capacities – at least not in a way that would make effectively performing the former inconceivable without an antecedent possession of the latter. One might even wonder if the insistence that there *must* be such a dependency does not merely assume the very uneasiness, we are trying to get clearer on… in which case, it would ultimately just be begging the question…

If one can only perform the work at hand on the basis of some antecedent and conceptually articulated understanding of shoemaking, then one cannot appreciate the

[27] Of course, one could hit a piece of leather on a shoe last with a hammer, without knowing or understanding anything about shoemaking. The point is that one could not do so in a way that would amount to shaping that piece of leather on the last. Not because one could not conceive of his action as such, but simply because one would not have the relevant skills. Another example might help here. Of course, one does not need to have learned how to play piano in order to randomly hit the keys of a piano. But at the same time, it seems, one *could not* play Rachmaninoff's Piano Sonata no. 2 in B flat minor, if one had not been taught how to play. The impossibility here is certainly a physical impossibility: without proper training, one simply could not have the required skill and dexterity to even hit the right keys at the right time. But I think we might want to go further and say that we cannot even really conceive of a situation in which someone who has never been taught how to play the piano – and might even never have touched a piano -, just sits down and randomly press keys in a way that ends up amounting to *playing Rachmaninoff's Piano Sonata no. 2*. Maybe it could be a scene in a movie, but, then I think we would end up objecting that despite what is claimed in the movie, the character *had to* have been taught piano… (In that way, the character playing the sonata might be similar to a pot seeing and hearing in a fairy tale or of a stone feeling pain (see *PI* §§282 and 284). At first sight, it seems we can get a grip on the idea – even though no pot or stone could *really* have sentience. But the more we look at it, the more we are forced to make concession to prevent the idea from collapsing on itself. It is, I think, what Wittgenstein is doing for example when he says of the pot that, in the fairytale, it can also talk…)

[28] This echoes a worry one might have when it comes to Wittgenstein's language game of the builder and the assistant. In Wittgenstein's initial description, we are encouraged to imagine the language game played by the builder and the assistant as "a complete language-game" (*PI* §2). But then one might wonder: could it really be a complete language-game (i.e. a language-game that could, in principle be played in isolation from any other such games)? That is, could two individuals really perform the kind of work described by Wittgenstein without having any broader conceptual or linguistic capacities?

hammer as fitting that work independently of that very understanding. And if one can only appreciate the hammer as fitting the work at hand on that condition, then it is difficult to see how the hammer could be given in perception in the way we have been suggesting. But what leads us to believe that the performance of the work at hand must be dependent upon an antecedent and conceptually articulated understanding of shoemaking? Maybe it is the impression that, it is only against the background of such an understanding that our perception of the hammer – as well as of the piece of leather and of the shoe last, etc. – could effectively guide the action of shaping the piece of leather on the shoe last. But then that is just another way of saying that the hammer could not be given to us in perception as having the relevant practical meaning, without the involvement of conceptual capacities. If this is the best we can come up with, then I fear we are running in circles.

Still, I think a concession can – and should be – made here. The shoemaker could not perceive the hammer as fitting the shaping of a piece of leather on a wooden last if he had no understanding of shoemaking. That much can and should be acknowledged. And this certainly forces us, more generally, to make room for understanding in our discussion of perceptual meanings. But, as I see it, it does not impose any constraint on what *form* this understanding must take – nor does it give it any logical or developmental priority on perception. As such, it is compatible with the suggestion that the understanding involved in the shoemaker's perception of the hammer does not consist in anything beyond that very perception. The understanding of shoemaking (e.g. of the steps and tools involved, etc.) might well manifest itself in the shoemaker's work – and it might even be a necessary condition for him to be performing that work at all. But such an understanding does not

have to be anything over and above the shoemaker's sensitivity and responsiveness to the perceptual meanings of the hammer, the shoe last, etc.[29]

(ii) Intentional Richness and Basic Behavioural Capabilities

This leads us to the second potential worry at the suggestion that meanings indigenous to perception can be intentionally rich: namely that if they were to exhibit such intentional richness, they would be too distant from our basic behavioural capabilities to be playing any genuine – and unmediated – guiding role. The issue with the suggestion that I could perceive the hammer as fitting a particular work would then have less to do with an alleged dependency of such practices on conceptual and rational capacities, and a lot more to do with the necessary grounding of complex actions on simple ones and that of cultural practices on biologically given motor skills. So, why might it be less contentious to suggest that the hammer can be given as *graspable* rather than as *fitting the shaping of a piece of leather on a wooden last?* Well, maybe because the action of grasping a hammer is a simple action – one that could not be reasonably decomposed in more basic actions. And not only is it a simple action, but it also manifests a motor skill that is afforded by our very bodily constitution. Given normal development, having hands like ours will simply translate into behavioural capabilities like that of grasping. In contrast, it seems the action of shaping a

[29] As I see it, this kind of understanding is precisely what Merleau-Ponty is trying to capture when he attributes to the typist "a knowledge located in the hands" (*un savoir qui est dans les mains*) (Merleau-Ponty, 1945, 179). It is a kind of knowledge or understanding that can be said to guide the typist's work, but which does not amount to the possession of conceptually articulated representations (of the keyboard, etc.). It is worth noting that, in this example, the hands are not merely performing the work; they are also a perceptual organ. There might be a temptation to read this passage as locating the understanding or knowledge in the very action performed 'by the hands'. But while Merleau-Ponty would recognise that the typist's understanding or knowledge *manifests* itself in the way he presses the key on the keyboard; I take it that, for him, this is only part of the story. As I take it, the understanding is best conceived as already 'guiding the hands' and so as best located in the perception the typist has of the keyboard.

piece of leather on a wooden last is a complex one: it involves various steps (e.g. grasping the hammer with one hand, holding the piece of leather on the wooden last, hitting the leather with the hammer, etc.), each of which can be conceived as an action in its own right. Moreover, the ability to so shape a piece of leather – inasmuch as it emerges in the context of a practice like shoemaking – is a cultural achievement. It depends on more than just our having hands and going through normal development. The simple actions that constitute the shaping of a piece of leather on a wooden last might well be actualisations of biologically given motor skills, but it does not follow that the ability to shape a piece of leather on a wooden last is, itself, such a skill.

Now, the thought goes, perception must certainly inform – and possibly even guide – simple actions.[30] Otherwise, it is difficult to see how we could make sense of our most basic success in practically engaging with the world (e.g. our managing to grasp objects, to avoid obstacles on our path, etc.). But perception does not have to separately inform and guide complex actions. Indeed, if such actions are not anything over and above the simple actions that constitute them[31], then it would be redundant to have perceptual meanings apt to play such role. Why would the hammer be given in perception as *fitting the shaping of the piece of leather on the last* if it is, for example, already given as *graspable*, as *a mean*

[30] I am distinguishing the idea of *informing* from that of *guiding* as follows: To say that perception informs certain actions is to say that perception delivers contents that are causally involved in the bringing about of said actions. In contrast, to say that perception *guides* certain actions is to say that this bringing about involves a receptivity to meanings and so a form of understanding. (We will come back to this distinction in section 4.3)

[31] To be clear, this is a claim about what *constitutes* a complex action – not about what the ability to perform it requires. As such, it is perfectly compatible with the suggestion that the performance of a complex action might require more than what is required to perform each of the simple actions that constitute it. Take again the example of playing Rachmaninoff's piano sonata no. 2. The ability to play the sonata certainly requires more than to have the motor skill to press piano keys. But one can accept that fact and still maintain that playing the sonata does not amount to more than pressing particular piano keys in a particular order and at a particular pace.

to hit, etc.? After all, the latter perceptual meanings appear quite sufficient to guide the simple actions of grasping the hammer, of hitting the leather with the hammer, etc. And if there's ultimately nothing more to shaping the piece of leather on a shoe last, then why posit perceptual meanings intentionally richer than those?

Furthermore, even if abilities emerging in the context of cultural practices are not themselves biologically given motor skills, it seems reasonable to assume that their possession is still *partly* dependent upon the possession of some such motor skills. After all, the shoemaker would not have the ability to shape the piece of leather on the last if he did not have the motor skills to grab things (in this case, the hammer), etc. Cultural practices might reorganise or repurpose our biologically given motor skills into new abilities, but there remains an inescapable dependency. The thought then is that just like the dependency of complex actions on simple ones, this second kind of dependency also speaks against the need for – and even plausibility of – intentionally rich perceptual meanings. If abilities involved in cultural practices like shoemaking are so closely related to biologically given motor skills, then why would they not be guided by the same perceptual meanings?

Now, Merleau-Ponty would certainly object to such a conclusion. Indeed, one of the central strands of the *Phenomenology* and of Merleau-Ponty's philosophy more generally is that we first and foremost inhabit – or situate ourselves in – a world that is inherently shaped by our cultural achievements. In perception, our surroundings are not usually given as merely biologically meaningful, but rather as culturally meaningful. To us, a hammer presents itself first as a tool fitted for various kind of works; it is only by abstraction, by actively ignoring the cultural context in which it is found, that it becomes a

mere *graspable something*.[32] In the same way, a pavement stretching along a street has perceptual meanings – dependent on there being cars, street signs, traffic laws, etc. – that goes beyond being simply a *walkable path*. For Merleau-Ponty, this is a general phenomenon. As he puts it: "I have around me roads, plantations, villages, streets, churches, utensils, a bell, a spoon, a pipe. Each of these objects carries in itself the mark of the human action to which it serves." (Merleau-Ponty, 1945, 404-405).[33] And this influences how the objects are given to me in perception. As Merleau-Ponty goes on to say: human actions "sediment themselves on the outside" (*se sédimentent* au dehors). That is, they come to be found *on* and *among* the objects around me.

As I read it, the image of sedimented actions suggests three things. First, the human actions to which objects serve – that is, their purposes – are somewhat *objectively* there. Of course, their presence depends on a series of cultural achievements and so, in some sense, depends on us.[34] But we do not merely project them onto objects. That's what Merleau-Ponty is trying to capture when he says that they come to have "the anonymous existence of things" (*l'existence anonyme des choses*) (Merleau-Ponty, 1945, 405). Second,

[32] In the same vein, it is only by abstracting altogether from our embodied situation, that a hammer comes to present itself not as fitting a work, or as graspable, but as a mere object with definite spatial dimensions, a definite weight, etc. (Of course, what makes a hammer graspable has to do with the diameter of the handle, its height and weight, etc. Likewise, what makes the hammer fit for the work of the shoemaker has to do, among other things, with it being graspable. The mistake, for Melreau-Ponty, is to then conclude that the hammer somewhat always has to be given to us in perception as having these definite properties that makes it graspable, etc.)

[33] It is worth noting Merleau-Ponty mentions both roads and streets in this passage. I do not think it is a useless contrast here. It might be if they were given to us in perception – maybe with pavements – as mere *walkable path*s. But I suspect, Merleau-Ponty would say that they are not, that roads and streets, though related, are very often given to us in different ways – i.e. contain 'the mark of different kinds of actions'. (Just like we might walk on a pavement, we might also walk on a road or a street, but, in each case, we might do so quite differently. Think of how relaxed one usually is when walking on a residential street, versus the hypervigilance one would most likely exhibit if forced to cross a motorway – irrespective of the presence or absence of cars in one's visual field. Then compare this with the lack of attention one usually has for incoming cars when walking on a pavement.)

[34] To borrow Brandom's expression, we might say that they are *creatures of ours* (see Brandom, 1998).

these purposes are there as something that can be – directly – perceived. If they are not projected onto objects, they are not hidden *under their surface* either. That is, they are not something that has to be inferred from some other perceptually available features. Third, inasmuch as they are *on* and *among* objects, their presence inescapably alters what it is we perceive when our attention is directed at the objects. They are not being given to us in perception as a mere addition to what would have been given, had they not been there.[35]

While enlightening in this respect, I think the image of sedimented actions nonetheless becomes misleading if we take it as implying that to see the 'real' objects we need to somewhat dust the sedimented actions off them. I think that's why Merleau-Ponty simultaneously talks of objects 'carrying in themselves' the mark of human actions. What might seem like a clumsy superposition of two conflicting images (i.e. one suggesting that human actions are somewhat *in* the *objects*, at their very core, the other that they are *on* the objects as something external to them) is, as I read it, rather a way to keep them in check, to prevent them from leading us astray.[36] In perception, the hammer *is* this graspable something, that fits the work of the shoemaker or that of the woodworker, etc. [37]

---

[35] Another interesting thing the image of sedimented actions suggests is that of a sustained historical – or at any rate temporally extended – process.

[36] As far as I can tell Merleau-Ponty does not *explicitly* reflect on his choice of images here, but it is worth noting that he often discusses perceptual meanings in similarly – apparently – conflicting terms. He'll say for example that the "meaning of a thing *lives in* the thing" (*le sense d'une chose habite cette chose*); that it "animates" the thing. But at the same time that it is not hidden "behind the appearances" (*derrière les apparences)*, behind what we see of the thing (Merleau-Ponty, 1945, 375).
In passing: I think we should see Merleau-Ponty here as trying to reclaim these various philosophical 'images' or 'pictures'. Wittgenstein warns us that pictures can hold us captive (see *PI* §115). But this captivity is not unavoidable. If we can be slaves to pictures, we can also make them work for us. Merleau-Ponty's strategy might be exemplifying precisely that.

[37] Think of dust on a bust. We might still perceive the facial features, even quite distinctly, but there might be no way to say what colour the bust is. So not only is the dust there to be seen, it also affects what we can perceive of the bust. If the layer of dust were to become thick enough, it might even become impossible to distinguish the contribution of the bust's shape to the overall contour of the thing as a whole. In the same way, certain objects are covered by so many layers of human actions that it becomes difficult – if not impossible – to distinctly *see* what lies underneath. We might well infer that the bare objects would look like this or that, if the layers of human actions were to be peeled away. But this is not what is really given to us in perception. In that sense, we could say that, in such cases, the relation between what is directly given in

Importantly, for Merleau-Ponty, what goes for perceptual contents, also goes for actions: the mistake is to think of culturally achieved abilities as biologically given motor skills *plus an extra something*. And so, to think of cultural achievements as merely reorganising or repurposing motor skills, without actually altering them. In contrast, the picture Merleau-Ponty is offering is one in which cultural achievements radically transmute these motor skills. Our interaction with cultural objects, our taking part in practices leads to the emergence of new capabilities, not just to the reorganisation and repurposing of old ones. It is on the ground of such considerations, for example, that he dismisses as misleading talks of our having a – culturally achieved –second nature *on top of* a biological nature (see Merleau-Ponty, 1942, 265 as well as 1945, 230-231).[38] For him, it is simply mistaken to assume that biologically given motor skills – or, as he often puts it "vital behaviours" – somewhat subsist untouched in the acquisition of new abilities (see Merleau-Ponty, 1942, 273). [39]

The argument was that, properly conceived, every action could either be appropriately guided by perceptual meanings closely aligned to our most basic behavioural capabilities or decomposed into actions that could. It rested on the assumption that biologically given motor skills were like building blocks for culturally shaped abilities. But

perception and what has to be inferred runs in the opposite direction to the assumed mereological dependency of the whole on its parts. (And that's why taken too far, the picture of sedimented actions becomes misleading: it suggests that, just like with the dust, the sedimented actions are wholly external to the object, and, in a way, 'hides it'.)

[38] It is along the same lines that he objects to talks of man as a "reasonable animal" (*animal raisonable*). The emergence of rational capacities does not leave intact a self-enclosed sphere of animal instincts (see Merleau-Ponty, 1942, 274).

[39] This stance might find its clearest expression in the *Structure of Behaviour*. But it is also implicit in crucial passages of the *Phenomenology.* For one, it is central to the claim that we orient ourselves towards a culturally shaped world (see Merleau-Ponty, 1945, 404 and 418 most notably). Such an orientation towards – or situatedness in – a world is as much a matter of what we can perceive of the world as it is of what we can do in the world.

Merleau-Ponty rejects precisely that too. Such skills or vital behaviours might bear witness to a certain mode of engagement with the world, but it is a mode of engagement that is largely superseded in the course of ordinary initiation into cultural practices. If there is a dependency of culturally shaped abilities on biologically given motor skills, it is best understood as a developmental one – rather than as a mereological one. And so, it is not so obvious after all that such a dependency speaks against the intentional richness of perceptual meanings – at least not by itself.

As for the dependency of complex actions on simple ones, I think there might also be reasons to doubt it has the implications initially assumed. Part of the issue is that, despite its apparent self-evidence, the distinction might prove difficult to fully flesh out independently of any consideration as to what is given to us in perception. After all, what can we appeal to in judging that an action is simple? On what ground can it be said that it cannot be decomposed further? Surely it cannot be a matter of it corresponding to a discrete bodily movement. Think again of the action of grasping the hammer. It ultimately involves a series of bodily movement (e.g. an initial extension of the arm, the positioning of the hand in a grip, the contraction of the fingers, etc.). And even these bodily movement can be further decomposed: the extension of the arm involves the extension at the elbow, but also a raising motion at the shoulder. The fingers are coordinated into a grip position, but this involves movements in each individual finger, etc. Do we want to rather say that *these* movements are simple actions – and so that actions like that of grasping a hammer are really complex ones? I think we probably would not. But let us suppose that we did. Even then, we might still be tempted to say that forming a grip with all fingers is easier to perform than forming a grip with all fingers, but one – and so that it is, in a sense, a 'simpler'

movement… Of course, it might not be simpler when we think of bodily motions in terms of geometrical transformations of body parts – and of underlying pattern of muscle activations. But it might be when we rather think of such bodily motions as manifesting more or less privileged modes of engagements with the world.

So, what makes the forming of a grip with all fingers more privileged *to us* compared to the forming of a grip with all but one finger? Well, if we accept Merleau-Ponty's view that our cultural achievements fundamentally shape the way we orient ourselves in – and engage with – the world, it seems wrongheaded to invoke here the idea that such a grip might be privileged because it is a biologically given motor skill.[40] So what else could make such a grip more privileged to us? I suspect, it is tempting to then invoke the content of our perception. Maybe it is a more privileged grip because the perception we have of the hammer is such that it can, by itself, guide the forming of such a grip – whereas the forming of a grip with all but one finger would require some further cognitive steps.[41] It is more privileged, because there would be a – more – direct pathway from how the hammer presents itself in perception to the forming of such a grip. The hammer would be given to us in perception not only as graspable, but as graspable *in that way*. And this would ultimately explain why the forming of such a grip is ultimately simpler. But then, if this is

[40] Another example might prove useful here. Take a teacup. In contrast with the hammer, the more privileged grip *for* us might well be with all but one finger (i.e. the little finger remaining relaxed). And surely there are times in which we might explain such a grip being more privileged by merely appealing to the size of the teacup's handle. If it is a more privileged grip, it might simply be because the handle is too small to accommodate all fingers. But, what if the handle is big enough to accommodate all fingers? Might it not be the case that the grip with all but one finger is still the more privileged *for us*? One plausible explanation – one to which Merleau-Ponty would, I think, be sympathetic – is that our engagement with teacup is shaped by a certain history of tea drinking and so by a number of cultural norms and habits that have 'sedimented themselves' on teacups.

[41] Although he does not quite express it in those terms, I think Merleau-Ponty hints at something along these lines when he discusses the relation between the idea of an action being simple and it being privileged (see Merleau-Ponty, 1942, 223).

how we flesh out the distinction between simple and complex actions, it can hardly speak against intentionally rich perceptual meanings, without already assuming that there cannot be such meanings! So conceived, the argument is circular. Of course, if simple actions are actions for which there is a direct perceptual pathway *and* there is no such pathway for actions like shaping a piece of leather on a wooden last, then such actions are not simple. But to assume that there is no such pathway is precisely to assume that meanings like *fitting the work of shaping a piece of leather on a wooden last* are not given in perception. This was supposed to be the conclusion of the argument, but it looks like it might also be a premise of it…

Now, I do not think that we have thus shown once and for all that there are intentionally rich perceptual meanings. However, I do think we have made a strong case against dismissing the suggestion outright. We have tried to give voice to two distinct worries, as charitably as possible. It might appear intuitively obvious that our conceptual and rational capacities have to play a role in the enjoyment of such meanings – in a way that is ultimately incompatible with the details of the suggestion. Or that such meanings are too distant from our basic behavioural capabilities to even be given in perception at all. However, despite our best attempts to flesh out these worries, they have either collapsed on themselves or proved question begging.[42]

[42] Here is another way to look at it: our argumentative strategy has been to show that our initial uneasiness towards intentionally rich perceptual meanings is not sustained by any decisive *a priori* argument. And if no such argument can be provided, then there is no reason to assume any upper limit to the intentional richness of perceptual meanings. Now, could it be argued that, while intentionally rich perceptual meanings are *conceivable*, they simply are not, *as a matter of fact*, the kind of things that are given to us in perception? Possibly. Although, it is not quite clear what kind of empirical evidence – if any – could really prove the proposal wrong… This echoes an issue we have already raised: while Merleau-Ponty's work is clearly empirically *informed*, it does not appear *grounded* on any particular empirical evidence. (Now, I do not pretend this distinction is all that obvious, nor that it would necessarily sustain further scrutiny, but exploring the issue would take us too far off our current concern.)

2.6 Motivation, Guidance and Objectivity

Why does this matter? Well given our current concerns, it matters for at least two reasons. First, we want to make sense of perceptual meanings being able to guide action in a way that is neither merely causal nor dependent on our possession of fully-fledged conceptual and rational capacities. Or, to come back to the notion of motivation: we want to make sense of objects being given in perception in such a way as to *motivate* certain actions. But for such talk of motivation to even get off the ground in cases like those of the shoemaker or the woodworker, we need to make room for rich enough perceptual meanings. If claims like "in the workshop, the hammer motivates the shaping of a piece of leather on a wooden last" are to be informative at all, then it has to be possible for the hammer to be given in perception in a way that could indeed guide the shaping of a piece of leather on a wooden last. Otherwise, it is not at all clear what such claims would be expressing – and a fortiori how they would be a genuine alternative to claims in terms of reasons. Indeed, if there are intentionally rich meanings, but they cannot be given in perception – or at least not given in perception in the way advanced by Merleau-Ponty –, then it seems claims about motivations might simply collapse with claims about reasons…

And this has to be a genuine concern for Merleau-Ponty. After all, he does maintain that, in ordinary circumstances, the best way to understand works like those of the shoemaker or the woodworker is in terms of motivation.[43] And, for him, the same goes for a great many kinds of actions. Merleau-Ponty sees motivations as the best way to capture our *default* mode of engagement with the world. And this goes from our dealing with

[43] See again Merleau-Ponty, 1945, 136.

'basic' obstacles and affordances (a rock on our path, a stretch of compacted soil as an alternative route, etc.) To our using tools like hammers or typewriters. All the way to our responding to erotic, social or political situations. These later cases, I think, are often missed. Many, sympathetic to Merleau-Ponty's proposal, will be happy to concede that there are perceptual meanings such that of being *an obstacle* or *a path* or *graspable* – and possibly even meanings such as *fitting a particular work* or *activity*. But I think many would nonetheless be hesitant in accepting that a partner's body can be given in perception as *saturated* with erotic meanings. In the same way, many would object that one cannot actually *perceive* a situation of social tensions and civil unrest as *revolutionary*. A 19th century, proletarian might well think of the situation as calling for a revolution. He can reason on this, conceive of himself and of the other protesters as proletarian, as unfairly poor, as exploited, but his material condition cannot actually be given to him in perception *as calling for a revolution*.[44]

[44] Dreyfus is an interesting exception. At times, his way of discussing Merleau-Ponty's views on perception and action may give the impression of such a blind spot. He often comes back to the same examples of a football player on a field, of a driver or of someone catching a ball, with a particular insistence on "skillful responses" and "coping", on getting a "maximal grip" on one's ever evolving environment (See for example Dreyfus, 2006, 136 and 138; 2017, 30-31 and Dreyfus and Taylor, 2015). Now, some of these examples, like that of the football player, are from Merleau-Ponty himself (see Merleau-Ponty, 1942). And I do not think Dreyfus, in discussing them, is unfaithful to the spirit in which Merleau-Ponty initially introduces them. My point is rather that, it may be tempting to see such examples of expertise as not depending on especially rich perceptual meanings.
But these are not the only cases of expertise Dreyfus ultimately frames in terms of "skillful responses", "coping", etc. That's also how he talks of ethical expertise – and more generally of any kind of "cultural" expertise. And these clearly involve a responsiveness to *intentionally rich* meanings (see for example Dreyfus, 2017, 32-33. See also Wrathall, 2017, 12).
As I see it, Dreyfus' discussion of the catching of a ball or the driving of a manual transmission car makes it clear that perceptual meanings may be of great simplicity – but are still precisely that: meanings. It suggests that the threshold for intentionality is much lower than the empiricist and intellectualist pictures both assume. In contrast, what we have tried to show – and what Dreyfus' discussion of cultural mastery or expertise makes salient – is that, once we acknowledge that there can be intentionality *without concepts*, then there is no *a priori* upper threshold.

Now, I think this is a mistake in both cases; it merely reasserts the kind of worries we have already dismissed. As it stands, we have no more ground to reject such perceptual meanings than we have to reject the perceptual meanings of being *graspable* or *fitting a particular work*. If we accept that, in a workshop, a hammer can motivate a particular work, we can – and should – take seriously Merleau-Ponty's suggestions that, in an erotic situation, the body of one's partner can motivate certain gestures or that, in a particular social situation, one's material condition can *motivate* revolutionary actions (see Merleau-Ponty, 1945, 193-195 as well as 117, 421-422 and 506-510).

If the aptness of the notion of motivation for such a wide array of cases matter to Merleau-Ponty, it also matters to us. Not so much because all these actions (from avoiding a rock on one's path, to joining a revolution) will necessarily count as cases of rule-following, but because they exemplify a wide array of situations in which we should expect to find cases of rule-following. We should not assume that the proletarian's sensibility to the revolutionary meaning of his material condition – or the shoemaker's sensibility to the practical meaning of the hammer - has to be analysed as the manifestation of 'more basic' episodes of sensibility coordinated by extra cognitive steps. And we should not expect one's sensibility to the 'meaning' of signposts, orders or pieces of a game to be any different…

Some might already be of the impression that the analogy with tools will prove the most informative (after all, might we not say that a signpost is given in perception as *calling for a particular response* in a similar way in which a hammer is given as *fitting a particular work*?). But I would be wary of hasty judgements here – as I would of broad

generalisations.[45] Hence the importance of showing that meanings indigenous to perception can not only be intentionally rich, but intentionally rich in different ways.[46]

The second reason why these considerations matter for our current concern is that they help us better situate the kind of objectivity we have tentatively fleshed out in the previous chapter. Merleau-Ponty sometimes talks of us projecting our practical intentions or projects unto the world. Or of us projecting our cultural or political situation around us (see Merleau-Ponty, 1945, 111 and 170). But at the same time – and as we have seen –, he does not want to then conclude that the kind of meanings such projections make salient are not really *there to be found*. Had it not been for such projections, we might have failed to notice many meanings, but this does not make these meanings, mere projections. What is needed then is a notion of objectivity that can reconcile both ideas and so make sense of us genuinely *perceiving* meanings, while recognising that the availability of these meanings is ultimately dependent upon our actual mode of engagement with the world. We might not perceive objects as *graspable* if grasping them had never even been a bodily possibility. That much is clear. But it does not then follow that being *graspable* is something we 'put on' these objects rather than 'read off' them.[47] As we have hinted at earlier, this is, for

[45] Relatedly, one could dispute the extent to which the analogy is really informative, without further qualification. After all, it seems we have to acknowledge that the purpose of tools, just like our relationship to them, is incredibly varied. This is something stressed by Wittgenstein (see *PP* §§11-17).
[46] We'll explore this idea in some more details in section 4.3.
[47] As I read it, I think Merleau-Ponty might even be willing to concede that the availability in perception of meanings such that of being *graspable* might be contemporary to our having practical intentions with that very content. So, were it not for our sometimes intending to grasp objects – for our projecting such practical intentions unto the world – we might never perceive any objects as *graspable*. But for Merleau-Ponty, this does not threaten the idea that the meaning is somewhat *in* the objects. At most, it shows that certain features of the world will only ever be available to us if we are appropriately attuned to them (I think this is what Merleau-Ponty is hinting at when he assimilates our reading off the meanings that "resides in" objects, to a sort of "dialogue" or "communion" with those very objects (see Merleau-Ponty, 1945, 375-376).

Merleau-Ponty, the problem of "existential modality" or of "transcendence" which concerns any meaning indigenous to perception (see Merleau-Ponty, 1945, 422-423).

To pretend to objectivity, the thought often goes, there must be a possibility of error. Something cannot be *objectively* there if the only criteria for it being there is that we think it is. It must be possible for us to be mistaken in our judgements about it. But there might also be something to be said in terms of perception: if something is objectively there, then it must be possible for us to *fail* to notice it. And so the fact that certain meanings are not always *given to us* in perception (e.g. that I might not see the hammer as *fitting the work at hand* or even see it as graspable), instead of speaking against the objectivity of such meanings, might, on the contrary be an argument for it.

As I see it, these considerations ultimately give us a blueprint for how to link three of the central insights we are trying to flesh out namely that (a) in ordinary circumstances, our responsiveness to signs – or to expressions of a rule – rest on our understanding of their meaning – or requirements, (b) such an understanding can be both non-conceptually articulated and given to us in perception (c) the meaning of signs – or the requirements of rules – are an objective matter.

2.7 A Layered Conception of our Engagement with the World?

So Merleau-Ponty is committed to the idea that we first and foremost inhabit a culturally shaped world. And, for him, this has implications both for what is there for us to perceive and for what is there for us to act on – or direct ourselves towards. Indigenous to perception – and action[48] –, there are, among other things meanings. And not only

[48] Initially, we contrasted meanings *indigenous* to perception to meanings that can only be enjoyed given the active involvement of fully-fledged conceptual and rational capacities. But given, Merleau-Ponty's insistence

meanings, but meanings dependent upon cultural achievements. At the same time, he is committed to the idea that our sensibility and responsiveness to these meanings do not ordinarily depend upon any kind of conceptually articulated and rationally structured understanding.[49] Taken together these two ideas seem to imply a rather loose relationship between our everyday engagement with the world and our conceptual and rational capacities. And this might seem rather problematic to some.

As we have seen, Merleau-Ponty rejects the suggestion that our cultural inheritance leaves unaltered a more 'basic' mode of engagement with the world, that could simply be unearthed from under it. For him, even our most 'natural' dealings with the world are transfigured by our initiation into institutions and practices, by our acquiring a culture. When it comes to our way of being in the world, the contribution of our biological constitution and our cultural inheritance do not superpose themselves as two distinct layers.[50] In being initiated into practices and institutions, one becomes sensitive to certain features (of objects), that would otherwise not manifest themselves. Objects cease to present themselves as mere affordances or obstacles and also become tools, objects fitted for certain works, etc. These further perceptual meanings aptly guide certain actions that would, otherwise, not have been within one's reach. But the perceptual meanings of fitting certain works are not built out of some biologically given meanings. And these new works are not just a reorganisation of biologically basic motor skills.

---

that the relation between perception and action does not have to be mediated by such capacities, it follows that meanings indigenous to perception can also reasonably be said to be indigenous to action.

[49] While Merleau-Ponty is committed to the idea that, *as a matter of fact* our sensibility and responsiveness to these meanings do not ordinarily depend upon such an understanding, much of what follows could be brought forward with the weaker commitment that they *do not have to* depend on such an understanding.

[50] On this point, I think there is a clear convergence with a strand in Wittgenstein's later writings that is encapsulated in the notions of form of life and of natural history and which ultimately implies something like the "mutual absorption of the natural and the social" (Cavell, 1989, 44).

But are we not creating a difficulty for ourselves in telling such a story independently of any reference to our conceptual and rational capacities – as, it seems, Merleau-Ponty enjoins us to? The worry is something like this: our conceptual and rational capacities – once they are actualised – must certainly shape some of our dealings with the world. Otherwise, they would be idle. And it seems, they must do so in a way that permeates perceptual states. Otherwise, it is not clear that they could relate such states to actions as reasons – and so genuinely guide how we interact with the world. But then, if our conceptual and rational capacities shape *some* of our dealings with the world, it seems, they must shape all – and so permeate all of our perceptual states

To deny this would risk creating a line of fracture either between two fundamental sets of capacities or between two modes of engagement with the world. Indeed, if our conceptual and rational capacities are contributing only to some perceptual states, it seems to follow either that we have self-standing perceptual – and behavioural – capacities to which conceptual and rational capacities only *sometimes* contribute or that we engage with the world in two radically different ways. Under one mode of engagement, our perceptual states are conceptually shaped and rationally linked to our actions; under the other, there is simply no such involvement of our conceptual and rational capacities.

The first alternative implies that it is possible, if only in principle, to isolate the contribution of our conceptual and rational capacities from some more fundamental perceptual – and behavioural – capacities. The risk is that, in conceiving the contribution of the former sets of capacities as separable from the latter, we create a further difficulty for ourselves: that of accounting for their interaction.[51] The second alternative faces a

[51] The issue finds various expressions in the literature. In McDowell's *Mind and World* – and in a Kantian spirit –, it is a matter of reconciling the sensibility and the understanding. And the only way this is possible,

related difficulty. Surely, there is a certain *unity* to our dealings with the world. So, if we indeed have two distinct modes of engagement, they must be *integrated* enough so that seamless transitions are possible from one to the other. But then, if these modes of engagement really are *fundamentally* distinct, what would make such a transition possible?[52]

The worry then is that, in positing perceptual meanings that are not conceptually articulated and in dissociating the ability of such meanings to guide action from any form of practical reasoning, Merleau-Ponty is denying the *overarching* influence of our conceptual and rational capacities on our way of being in the world and, in doing so, is either denying them any influence *at all* or is succumbing to a problematic layered conception of our engagement with the world.

Now, I do not think Merleau-Ponty is blind to such concern – nor completely ill-equipped to address it. That being said, I do not think he has a fully worked-out answer either… Some aspects of his position remain somewhat programmatic. As I read him, Merleau-Ponty is indeed denying the overarching influence of our conceptual and rational capacities on our dealings with the world. But he does not accept the conclusion of the argument we have just sketched. How so? Well, the argument is especially gripping, it seems, if we take our perceptual states to be *mere causal* happenings and our conscious life to be largely a matter of exercising our conceptual and rational capacities, of navigating a

for McDowell, is to deny that the sensibility can work in isolation from the understanding. I think there is, on this point, a clear convergence with Conant's reading of Kant as critical of a "layer cake conception of human mindedness" (see Conant 2017).

[52] As I see it, these issues are made especially salient by Dreyfus's picture of a "conceptual upper floor" and a "ground floor" of "embodied coping" (see Dreyfus, 2005, 47). One may even wonder if, in his own picture, Dreyfus does not ultimately put too much distance between the upper and ground floors to ever be in a position to convincingly solve these issues of integration. (As I read it, this is precisely what underlies his dispute with Romdenh-Romluc (see Romdenh-Romluc, 2007 and Dreyfus, 2007).)

space of reasons. If that is the picture we start with, then of course, the question of how perceptual states can be situated within the space of reasons – or to put it differently: the question of how navigating that space can be involved in perceptual states informing or guiding actions – becomes quite perplexing. We are stuck, it seems, with two irreconcilable orders: that of reasons and that of mere causes. Our conceptual and rationality capacities could only 'make a difference' in our dealings with the world if our perceptual states were somewhat permeated by these capacities. It is only at that condition that they would be *reasons* rather than *mere causal happenings* – and so could be properly situated within the space of reasons.[53]

But, in introducing the notion of motivation, Merleau-Ponty precisely intends to escape this very picture of our dealings with the world and of our conscious life – and so, to sidestep many of the problems that comes with it. For Merleau-Ponty, motivations are not situated on either side of the divide between the order of reasons and that of mere causes.[54] But, at the same time and crucially, they are not irreconcilable with either of these sides. Rather, when it comes to our dealings with the world, mere causes should be seen as an impoverished take on what motivations really achieve. Whereas, reasons should be seen as crystallising motivations, rather than as superposing themselves on them, as a distinct and irreconcilable order. Hence, for example, Merleau-Ponty's insistence that the intentionality we find in motivations – and, as such, in all the meanings indigenous to perception – is a "condition of possibility" of the intentionality we ordinarily ascribe to

[53] I take this to be very much in line with the kind of argument McDowell invokes against conceptions, such as Davidson's, which situate perceptual states outside the boundary of the 'space of reasons' – and so, outside the sphere of influence of our conceptual and rational capacities (see McDowell, 1996, lecture I).
[54] Or as McDowell might put it: between the space of reasons and the realm of law (see McDowell, 1996, lectures IV and V).

reasons – that is, to conceptually articulated and rationally structured items of consciousness (see especially Merleau-Ponty, 1945, 492-493 and 496).

The issues of interaction and integration we have raised are especially gripping if we think of perceptual states as *mere causal happenings* and contrast them with our essentially *intentional* conscious life. And if the only way for our conscious life to be intentional is for it to be shaped by our conceptual and rational capacities, then of course, these issues appear inescapable once we deny the overarching influence of these capacities on our perceptual states. But Merleau-Ponty denies both premises. He does not consider perceptual states to be mere causal happenings; they are for him intrinsically intentional. And he rejects the idea that the intentional character of our conscious life is dependent upon our conceptual and rational capacities.

Of course, this implies that if we are to follow Merleau-Ponty we need to rethink what it is for something to have intentional content – in such a way that it ceases to be equivalent to it being conceptually articulated. Or, as we have framed it before: we need to make sense of our practical engagement with the world as revealing a genuine kind of understanding, but an understanding we can arrive at without any act of subsumption under a concept.[55] One possible way to do it – which I take to be roughly Merleau-Ponty's – might be this: a perceptual state – or more generally, a conscious state – has intentional content to the extent that its enjoyment is the enjoyment of an *internal relation* (see especially Merleau-Ponty, 1945, 51 and 76). Take for example my perception of the hammer *as graspable*. It has intentional content, not because of any representational character of the perception – or because of the involvement of any concept –, but because,

[55] See for example Merleau-Ponty, 1945, 179-180 and 375

in being so given in perception, the hammer is being given as *internally related* to something else. In this case, the action of grasping (see Merleau-Ponty, 1945, 76 and 258). And, ultimately, appreciating such an internal relation *is* in itself an act of understanding.[56]

Let us assume, for a moment that this fits the bill and that once fleshed out in some more details, this will help us properly account for the intentionality of perceptual meanings – and as such of motivations. It quickly raises a question: if our conceptual and rational capacities are not the (sole) guarantor of intentionality – and of understanding – anymore, what role is left to them? Now, Merleau-Ponty is fully aware of this and acknowledges it early on in the *Phenomenology*, but it is not clear that he ever offers any satisfying answer.[57] It does not take much then for the worry of idleness to reemerge. The influence of our conceptual and rational capacities might not be overarching, but it seems they still have to play some role in shaping some of our dealings with the world. Otherwise, how could we even make sense of us having such capacities? And so, any account of our engagement with the world that appears at pain to acknowledge any role to them has to be suspicious.

Ultimately, I think this is a worry that Merleau-Ponty cannot simply dismiss – and might not satisfyingly answer. His insistence that our conceptual and rational capacities are not the primary mean by which we understand the world and ultimately orient ourselves towards it certainly contribute to making this worry less 'worrying'. After all, if such capacities are ultimately peripheral to our way of being in the world, then of course less

[56] More on this later (see infra chapter 3 and section 4.3).
[57] Merleau-Ponty puts it as follows: if we take seriously talks of motivations, "the [kantian] understanding also needs to be redefined, since the general function of liaison that kantianism ascribes to it is now common to all intentional life and so does not suffice anymore in circumscribing it" (Merleau-Ponty, 1945, 79).

appears to be at stake. For one, any uncertainty about their role appears a lot less threatening to the overall intelligibility of our dealings with the world.

That being said, it does not make the worry wholly disappear. Rather, it simply makes it easier to push it back. To conclude this section, I want to briefly consider one way in which, I think, the worry might be more satisfyingly addressed. I think it is largely compatible with Merleau-Ponty's picture, but I do not think he ever explicitly endorses the proposal.

What is needed is a way to acknowledge that our conceptual and rational capacities are dependent – even parasitic – on a more 'fundamental', and in this case intrinsically practical, mode of engagement. But to do so without either making these capacities wholly superfluous to our actual dealings with the world or wholly separable from that practical mode of engagement. One way to open up the conceptual space so as to accommodate both constraints would be to distinguish the actualisation of our conceptual and rational capacities from their actual involvement in any particular case. An example will be helpful here. Let us come back to our shoemaker and let us imagine that he does master the concept of *hammer* – and relatedly that he is able to navigate the broader network of concepts in which it is ordinarily situated. The possession of such a conceptual mastery might very well enable him to perceive more objects as *fitting the work of shaping a piece of leather on a wooden last*. That is, his being able to reason on what makes a hammer *a hammer*, on its relevant features, etc., might ultimately enable him to see object he would otherwise not have seen as *fitting such a work*. So, in that sense, the actualisation of his conceptual and rational capacities might, in some cases, have an effect on his dealings with the world. But the effect is one of *disclosing* perceptual meanings that would otherwise have been missed,

not necessarily of *shaping* these meanings. Or, to put it differently: were it not for his conceptual mastery, he might have failed to notice that certain objects other than – let us say – *his* hammer are also fitting the work at hand. But this does not then imply that when such a practical meaning is given to him in perception, it always involves or depends on an actual cognitive act of subsumption under a concept.

As such, the actualisation of our conceptual and rational capacities might have an effect on our overall way of being in the world – and even an effect that cannot be easily isolated. If they are parasitic on a more 'fundamental' mode of engagement, the suggestion then is that they also feed back into it.[58] They contribute something that was not already there. But this can be conceded while still denying that they permeate or are actively involved in every perceptual – and more generally conscious – state. Even if the shoemaker masters the concept of *hammer*, it does not follow that his hammer is always given to him in perception as subsumed under that concept.[59]

In response to Dreyfus, who accused him of succumbing to the "myth of the mental", McDowell has suggested that it is rather Dreyfus – and with him, many of his "phenomenologist heroes" – that succumb to a myth: that of the "Disembodied Intellect" (McDowell, 2013c, 322. See also Dreyfus, 2005). For McDowell the myth amounts to a conception of our conceptual and rational capacities as completely detached from our ordinary practical engagement with the world (see McDowell, 2013c, 322). And it is a

[58] Just like with the biological and the cultural, it would thus ultimately be misleading to talk of superposed layers.

[59] Now this is still a very programmatic proposal. A more complete story would involve some considerations as to *why* the actualisation of our conceptual and rational capacities can disclose – or make salient – meanings we would have otherwise missed. As I see it, such a story would probably revolve around the idea that the actualisation of these capacities enables us to 'crystallise' perceptual meanings in such a way that abstract manipulations become easier.

myth because, as appealing as it might initially appear, it ultimately proves untenable. Our proposal would certainly avoid that myth.

2.8 The Unfinishedness of Phenomenology

To conclude this section, I want to briefly reflect on Merleau-Ponty's conception of phenomenology and its possible achievements. More specifically, and in light of our discussion in the previous chapter, I want to look at the extent to which what Merleau-Ponty says about phenomenology might align with Wittgenstein's quietism.[60] Some affinities have previously been suggested between both views – and even more generally between the phenomenological tradition and a certain way of inheriting Wittgenstein's philosophy or at least of doing philosophy after Wittgenstein (see Glendinning, 2007, 5-6 and 24-28). Here is one way to further flesh out these affinities.

In the preface to the *Phenomenology*, Merleau-Ponty reflects on phenomenology, its historical development and its underlying methodological commitments. At one point, he claims that "the unfinishedness of phenomenology and its inchoate appearance are not a sign of failure; they were inevitable because phenomenology's task is to reveal the mystery of the world and the mystery of reason." (Merleau-Ponty, 1945, 21-22). Now, it may be tempting to read this passage as suggesting that the state of unfinishedness in which phenomenology has found itself – and maybe still find itself – is only temporary. And so that phenomenology will, one day, reach a point where it has said everything it has to say.

[60] On Wittgenstein's quietism, see supra 1.5-1.7

The unfinishedness of phenomenology, on that reading, would not be a constitutive feature of it, but a contingent fact of its history.[61]

But there is another way to read this passage: If phenomenology is unfinished, it is because of the nature of its aim – and ultimately because "the mystery of the world and the mystery of reason" resist a definitive elucidation. Now, Merleau-Ponty does not seem to suggest that phenomenology's aim *has* changed – one may even wonder if it *could* change without phenomenology becoming something else entirely. And so, as long as its aim is "to reveal the mystery of the world and the mystery of reason", phenomenology will remain unfinished.

Now, on this second reading, the unfinishedness of phenomenology becomes a constitutive feature of it. And this, obviously, seems to echo our suggestion about Wittgenstein's own conception of philosophy's task: as something that is without end (see especially *Z* §447). In both cases, there seems to be a keen awareness that this may appear disappointing. In Wittgenstein, the disappointment is voiced by a fictive interlocutor: "But in that case [if this is indeed how we do philosophy] we never get to the end of our work!". In Merleau-Ponty, it is not so much voiced as it is pre-empted by the reassurance that the unfinishedness of phenomenology is not a mark of its "failure".

If Merleau-Ponty is interested – and considers phenomenology to be interested – in the "essence of perception" (Merleau-Ponty, 1945, 17) then, it is not, I think, as an infinite longitudinal strip that could be grasped at once. It is rather as an open-ended collection of

[61] The temptation is, I think, reinforced by a possible mistake in both English translations. Landes translates the passage as follows: "The unfinished nature of phenomenology and the inchoate style in which it proceeds are not the sign of failure; they were inevitable because phenomenology's task *was* to reveal the mystery of the world and the mystery of reason." (my *emphasis*.) The past tense seems to me to be mistaken here (the French original reads: "[…]parce que la phénoménologie *a* pour tâche […]"). It strengthens the impression that what contributed to the unfinishedness or inchoate appearance of phenomenology is *in its past*. Smith's translation of that passage is distinct from Landes on a few minor points, but not on the use of the past tense.

variously related phenomena, as a potentially endless series of finite cross-strips that require our constantly renewed “attention” and “astonishment” (Merleau-Ponty, 1945, 22).

Chapter 3

Wittgenstein on Intentionality and Internal Relations

3.1 Introduction

Our issue now is the following: if we are to take seriously Merleau-Ponty's notion of motivation, we need to make sense of how intentional states (i.e. states with intentional content) can be internally related to features of the world even when they appear entirely dissimilar to each other.

I think a useful way to illustrate this issue of how to conceive internal relations involving intentional states is to look back at Wittgenstein's *Tractatus*, as well as at *Philosophical Investigations*. A popular narrative is that the shift from the *Tractatus* to *Philosophical Investigations* consists in the recognition that such relations cannot hold *de re* but only *de dicto* (see Hacker 1996, 2013a and 2013b and Crane 2010 and 2014). I think this narrative is mistaken. Not only does it misrepresent both the *Tractatus* and *Philosophical Investigations*: it also makes it impossible to articulate what I take to be a more satisfying stance on the general issue.

3.2 Intentionality in the *Tractatus*

In the *Tractatus*, Wittgenstein appears to be conceiving internal relations as grounded on essential properties.[1] But many of the points he is making can be made in

[1] This claim might be disputed. Crane, for example, is quite insistent that Wittgenstein's understanding of the notion of internal relation does not align with contemporary views – views according to which internal

conceiving such relations as grounded on merely intrinsic properties. One central insight of the *Tractatus* is that a picture can only be a picture of something with which it shares a common *form* (see *TLP* 2.171). When the idea is first expressed, it might seem that it is too strict a requirement. After all, is it not the case that something spatial – like a string of letters –, for example, can be a picture of something that is not spatial at all? However, this worry quickly disappears once it is established that all pictures – and more generally all states of affairs – ultimately have a logical form. And so, the requirement of shared form is not after all too strict as any picture – whether a spatial one or a "coloured" one, etc. – is also a logical picture. (see *TLP* 2.182).

Because of this early concession, one might get the impression that the issue we are interested in is largely absent from the *Tractatus* – not in the sense that Wittgenstein would have been blind to it, but in the sense that his early conception of meaning would be largely immune to such concerns. And this impression might be strengthened by the following

---

relations are grounded on so-called non-relational properties (see Crane, 2010 90-91. For a discussion of such contemporary views and on the metaphysics of internal relations more generally see Heil 2009 and 2012). According to Crane, Wittgenstein considers internal relations to be essential to their relata (see Crane, 2010, 90 and 92. On this point he considers himself to be following Beaney 2006). And, of course, this would imply that internal relations themselves would be essential to their objects rather than being grounded on some non-relational properties. That is, Crane attributes to Wittgenstein the view that internal relations are ultimately not ontologically superfluous.

I contend this is not Wittgenstein's position. To motivate his reading, Crane invokes the following passage of the *Tractatus*: "This shade of blue and that one stand, eo ipso, in the internal relation of lighter to darker. It is unthinkable that these two objects should not stand in this relation" (*TLP* 4.123). But there is clearly a way to interpret this passage that brings it very much in line with contemporary views. As I see it, Wittgenstein's point, in this passage, is that the two shades of blue *necessarily* stand in such a relation – and that it is necessary that they so stand *because* their being the shade they are is essential to them. Being *this* shade of blue is, for Wittgenstein, an "internal property" – that is, a property such that it is "unthinkable that its object should not possess it" (*TLP* 4.123). So, the internal *relation* of lighter to darker holds on the ground of the internal *properties* of the objects. And it is these properties that are essential to their objects – not the internal relation they ultimately ground. (Note here that, in Wittgenstein's example, the objects are the shades of blue themselves. And the internal properties ascribed to them is being *the shade of blue they are*. So it is, in a way, undisputable that *these* are indeed essential to their objects.)

Crane's reading is also at odd with Wittgenstein's discussion of the internal relations that hold between "the structure of propositions" (see *TLP* 5.2-5232). There, it becomes quite obvious that the internal relations are *grounded* on the structure of propositions themselves; they are nothing over and above it.

observation: the closest we come in the *Tractatus* to a sense of wonder as to how an internal relation like a relation of picturing could hold between two states of affairs that appear dissimilar is in 4.011. There Wittgenstein does acknowledge that "at first sight a proposition – one set out on the printed page, for example – does not seem to be a picture of the reality with which it is concerned."[2] However, he dismisses this worry as quickly as he raises it. As he goes on to say: "but neither do written notes seem at first sight to be a picture of a piece of music […] And yet these sign-languages prove to be pictures […] of what they represent" (*TLP* 4.011).

The underlying thought here is that the appearances simply prove to be misleading. Despite seeming entirely dissimilar, they have in common a shared logical form. And that's why relations of depiction can hold between them. If every state of affairs has a logical form, then the requirement that relations of picturing be grounded on such shared form certainly will not be too strict. But this is not the end of the matter, as the requirement then appears to be too broad in at least two ways.

First, under a rather common view, relations of picturing are fundamentally *asymmetrical*. That is to say if a state of affairs p is a picture of a state of affairs q, q is not also *ipso facto* a picture of p. For example, a musical script might be a picture of a symphony, but that does not make the converse true. Indeed, it seems rather counterintuitive to think of the symphony as being a picture or representation of the musical

[2] In talking of a proposition, Wittgenstein here already assumes that we are talking of a propositional sign with a sense – i.e. of a propositional sign that is in a determinate pictorial representation with a possible state of affairs. But the worry could certainly be rephrased in such a way that we do not prejudge whether or not the propositional sign – in this case, a series of letters on the printed page – have a sense. The worry can be expressed as follows: "how can a series of printed letters ever be a picture of a state of affair that does not resemble at all the said letters?". And the answer of the *Tractatus* is of course that all states of affairs – which include propositional signs - already exhibit what is essential to depiction, namely a logical form. This is why it is no more problematic for strings of letters to depict state of affairs than it is for hieroglyphic scripts (see *TLP* 4.016).

script – even if the latter is a faithful transcription of the former. Still, I think such a view is not obviously easy to articulate within the framework of the *Tractatus*. Indeed, the issue emerges from the recognition that pictures are states of affairs[3] and the suggestion that the possibility of depiction depends *only* on shared logical form. If any picture is a state of affairs, then it is also something that can be *depicted* by any other state of affairs with which it shares its logical form. They all have the necessary "formal properties" (see *TLP* 4.122). And so, within the framework of the *Tractatus* it seems nothing would prevent the symphony to be a picture of the musical script. As a matter of fact, Wittgenstein appears to concede that much when he says of the gramophone record, the symphony and the musical script that they "all stand to one another in the same internal relation of *depicting* […]" (*TLP* 4.014, my *emphasis*). There is no fundamental ontological difference between pictures and state of affairs depicted, in the *Tractatus*. And since the capacity of the former to depict the latter rest precisely on them having a shared logical form, then states of affairs depicted seem to have all that is needed to also be depictions of their pictures.

This is certainly one way in which Wittgenstein's requirement on relations of depiction might appear too broad. In itself, it cannot maintain the asymmetry of such relations.[4] Something more is needed to distinguish states of affairs that have the formal properties to be pictures, but ultimately are not, and states of affairs that not only have such properties, but also *are* pictures.

We must steer clear of a potential confusion here: it is largely accepted that the *Tractatus* – or at least the picture theory which finds expression in the *Tractatus* – is

[3] Wittgenstein says more specifically that they are facts (see *TLP* 2.121). But since facts are existing states of affairs (see *TLP* 2), there is no impropriety in saying that pictures are *states of affairs*.
[4] To be sure, many internal relations are asymmetrical. Just take the relation of being a darker shade of blue (see *TLP* 4.123)

primarily concerned with the *possibility* of representation (see Crane, 2010, 90; Hacker, 2013a, 50 and 2013b, 175). And so, one might be tempted to conclude that this possibility is precisely not disputed when it comes to the asymmetry of relations of depiction. The thought is something like this: the asymmetry raises an interesting question, but one that ultimately has little bearing on our current concern. It forces us to ask: what makes *this* state of affairs rather than *that* one a picture? But in doing so, it implicitly acknowledges that the possibility of pictures is not under threat.

I contend that this is a mistake. The issue of what makes representation possible is the issue of what can ground relations of depiction. That is, relations that hold between a picture and what it depicts. Of course, that every state of affairs has logical form makes it so that every state of affairs has the formal properties necessary to be a picture. And so, in that sense it gives us the assurance that every state of affairs *can* be a picture. But this is not the possibility we are trying to secure. What we are trying to get a grip on – and, as I take it, what Wittgenstein is trying to get a grip on in the *Tractatus* – is rather the possibility of *this* state of affairs being a picture of *that* state of affairs. And the asymmetry of relations of depiction is of importance here – precisely because it shows that shared logical form, in itself, is not sufficient to ground such relations.

So what more is needed to preserve the asymmetry of relations of depiction? One suggestion might be that the distinction between a state of affairs having the formal properties to b*e* a picture and it actually being a picture boils down to the use we make of it. So, the symphony is not a picture of the musical score – not because something about the logic of depiction makes it impossible, but simply because we do not use the symphony

in that way. We do not use it as a sign or picture and that is why it is not one. Now, as I see it, such a move would not be out of place in the *Tractatus*.

It is precisely the kind of considerations Wittgenstein appeals to when he introduces the notion of symbol, in contrast with that of sign. In the *Tractatus*, a symbol is a sign *with a sense* – whereas a sign is merely "what can be perceived of a symbol" (*TLP* 3.32). So, if the sign is a propositional sign, the corresponding symbol is a proposition.[5] The issue, of course, is that a same propositional sign can end up being part of different symbols – and so might end up having different senses or meanings. So how can we ever recognise from the – propositional – sign, the right symbol? In the *Tractatus*, the answer appears to be: we cannot from the – propositional – sign alone, we also need to take into account how it is "used", its "logico-syntactical employment" (*TLP* 3.326-3.327). Of course, this issue is somewhat distinct from ours, but it shows that the move of appealing to our actual "use" of "employment" of a sign is available within the framework of the *Tractatus*.[6]

Notwithstanding these considerations about the asymmetry of relations of depictions, the example of the gramophone record, the musical script and the symphony reveals a second way in which the requirement of shared form is simply too broad. As we have noted, Wittgenstein himself repeatedly insists that the relation of picturing between thought and reality depends upon them sharing a *logical form*. And commentators, like Hacker, have insisted that this is significant inasmuch as it enables Wittgenstein to claim that thoughts and propositions are pictures of *possible* states of affairs – thus giving himself enough

[5] As such the distinction between a sign and a symbol is not the distinction between a physical sign and its meaning, but between the physical sign, by itself – the latter Wittgenstein might say standing 'dead' (see *PI* §§432 and 454) –, and the physical sign as having a meaning

[6] The issue raised in *TLP* 3.32s already assumes that what we are dealing with are – propositional – signs and so, ultimately, pictures. Whereas our issue is precisely to determine what makes certain states of affairs pictures to begin with.

conceptual room to reconcile the idea that the relation between thought and reality is one of *picturing* with the possibility of false propositions (see most notably Hacker 1996, 2013a and 2013b). But, as commentators like Hacker also concede, this cannot be the whole story. That a proposition has a logical form might guarantee that it has the necessary formal properties to be a picture. The insistence that logical form is what any thought or proposition must have in order to be able to represent reality in general – given that logical form is indeed the form of reality (see *TLP* 2.18) – captures precisely that. But is shared logical form sufficient to ground a relation of depiction between *this* proposition and *that* state of affairs?[7] It seems not: if there is no further requirement, then any state of affairs can be, in principle, a picture of any other state of affairs…

But in various passages, Wittgenstein seems to anticipate this and to suggest a further requirement: the isomorphism must not only be one of form, it must also be one of structure. If there is – or at least can be – a relation of picturing between the musical script, the symphony and the gramophone record, it is not simply because they all have a logical form, it is also because they are constructed according to "a common logical pattern" (*TLP* 4.014). Every state of affairs has a logical form and as such, every state of affairs has a structure – i.e. a determinate way of arranging their constituting elements.[8] But not every state of affairs has *the same structure*.

The relation of depiction between the musical script, the symphony and the gramophone record depends on them exhibiting the same structure in such a way that a one-to-one correspondence can be established between their respective constituting

[7] This is very much how Hacker frames the issue (see Hacker 1996, 23 and 2013a, 51). In contrast, Crane remains silent on this, and appears content to insist on the idea of a shared logical form (see Crane, 2010, 92).
[8] This is how, I think, we should understand the claim that "form is the possibility of structure" (*TLP* 2.033).

elements. This implies the idea that the state of affairs involved in relations of depiction must have the same "logical multiplicity" – they must be composed of the same number of elements (see Bank 2007). But it involves more than that: inasmuch as they have a structure, states of affairs are not just sets of elements, they are an arrangement of such elements *in a determinate way* (see *TLP* 2.031).[9]

Now, does an appeal to shared structure – rather than merely to shared form – solve our issue? The answer seems to be no. Just take the idea of logical multiplicity. What makes it so that *this* state of affairs has the same logical multiplicity as *that* one? We have said: they have the same number of constitutive elements. But how are we to decide how many elements a state of affairs is composed of? The metaphysics of the *Tractatus* seems to offer a pretty straightforward answer: states of affairs are a concatenation of objects. So, it is just a matter of getting to the objects, to these "simple" and "unalterable" elements (see *TLP* 2.02;2.027). It would then just be a matter of comparing the elements that compose the state of affairs depicted to the elements that composed the state of affairs that pictures it. Or to put it differently: it would just be a matter of relating to each object, the name that stands for it (see *TLP* 3.21).

This would imply that the logical multiplicity of a picture would be grounded on it having the formal properties it has – and so on it being not only a state of affairs, but a state

[9] Interestingly, Merleau-Ponty discusses a similar issue in the *Structure of Behaviour* while appealing to the same example as Wittgenstein's. Merleau-Ponty wonders how the melody, the musical script and the series of gesture of the musician can stand in a relation of "expression" with each other and so "communicate internally". He ultimately voices the following answer: if such internal relations hold between the melody, the musical script and the series of gestures of the musician, it is because the three "participate in a common *structure*, have in common an identical nexus of meaning." (Merleau-Ponty, 1942, 185). He even goes to say: "The true sign represents what it signifies […] inasmuch as it is linked to the other signs in the same as the signified object is linked to the other objects" (Merleau-Ponty, 1942, 185). We cannot go into details here, but I take such a view to be already in tension with other elements of the *Structure of Behaviour* – and to have been altogether abandoned by the time of the *Phenomenology*.

of affairs *with that determinate structure*. But is it really plausible? Is the logical multiplicity of a picture really essential to it? Is there really something about the formal properties of a written sign such that it could not be of any other logical multiplicity – that it could not be analysed as composed of a different number of "simple" elements? It seems not.[10] And this is one way in which we might want to say that signs are 'conventional' or 'arbitrary'.

The issue then, is that there is a tension between the suggestion that relations of depiction are grounded *solely* on the formal properties of states of affairs – and the intuition that, *qua* signs, states of affairs do not really have such formal properties. The former suggestion appears to be forced on us by the picture theory and its attending metaphysics. The later intuition is one way to voice the – apparently undeniable – thought that there is something 'conventional' or 'arbitrary' about signs. Taken, in themselves, signs simply do not have *a determinate structure* – and so, do not have the formal properties on which relations of depiction allegedly *must* rest. The issue then is not just that the requirement of shared structure – like that of shared form – might be too broad. The worry now is that the requirement ultimately appears self-defeating.

Interestingly this intuition that signs are, in a way 'conventional' or 'arbitrary' finds expression in the *Tractatus* (see *TLP* 3.322 and 3.342). It seems to underly much of the distinction Wittgenstein draws between symbols and signs. As such, the issue we are facing here appears to be ultimately internal to the *Tractatus*, even if it is not explicitly acknowledged. At one point, Wittgenstein claims the following: "so what is essential in a

[10] This very much seems like the conclusion Wittgenstein arrives at in his later writings (see especially *PI* §48), but my suggestion is that, even though it is not clearly stated there, such a conclusion is very much in the spirit of some concessions Wittgenstein already makes in the *Tractatus* regarding the arbitrariness of signs.

proposition is what all propositions that can express the same sense have in common" (*TLP* 3.341). But, remarkably, it is quickly followed by the recognition of the arbitrariness of our notations – and so of our propositions – or rather propositional signs. And so, this raises the question: once they have been stripped away of all their arbitrary features, is there even something left, common to all these signs? It would have to be their structure, their common logical pattern. But at the same time that very structure also seems, to some extent, arbitrary…

What then? How can Wittgenstein address this issue within the framework of the *Tractatus*? One proposal, which is I take is pretty much Hacker's, is to appeal to thoughts. Propositional signs might well be arbitrary in the way we have highlighted, but, in the *Tractatus*, thoughts are not. They have a determinate structure and as such a fixed logical multiplicity just like the states of affairs they depict. The issue we have raised can thus be circumvented, the proposal goes, by recognising that the pictoriality of propositional signs is derived from – or parasitic on – the pictoriality of thoughts. For Hacker, this is very much Wittgenstein's stance in the *Tractatus*, even though it ultimately proves to be untenable. In the *Tractatus*, the picture theory would be saved from the issue we have highlighted. But the price to pay would be a commitment to an untenable conception of thoughts (see Hacker, 1996, 23-24 and 2013a). For Hacker, the picture theory, as it is articulated in the *Tractatus*, is ultimately coherent, but it rests on a mistake – and it is a mistake that only becomes manifest in light of the kind of considerations Wittgenstein raise much later…

Now I think this is mistaken, for two reasons. First, I do not think the relation between thoughts and propositions – or propositional signs – is as Hacker claims it is. Indeed, in the *Tractatus*, Wittgenstein defines a thought as a propositional sign with a sense (see *TLP* 3.5-

4).[11] So then how could the pictoriality of propositional signs be dependent upon the pictoriality of thoughts? Hacker's proposal rests on the assumption that, in the *Tractatus*, thoughts are unlike any other pictures – and as such unlike propositional signs – that their pictoriality is grounded only on their essential – and so intrinsic – properties. But if they are indeed unlike any other picture, why are they defined by appealing precisely to other pictures? Pictures for which the issue we have raised remains whole.[12] Second, even if we were to concede to Hacker that there is indeed such an appeal to thoughts in the *Tractatus*, it would prove unsatisfying *by the Tractatus*'s own light! Think again of Wittgenstein's claim that what is essential to a proposition – or rather a propositional sign – is what all propositional signs that can express the same sense have in common. In light of Hacker's proposal, we would be tempted to say: what is common to them is that they express the same thought! But then, if a thought is a propositional sign *with a sense*, we are running circles. What is common to all signs that can express the same sense would be that they are expressing the same sense…

My own proposal is that, within the framework of the *Tractatus*, it is rather the notion of use or of employment that can ultimately be invoked in response to the issue we have highlighted. The tension between the requirement of isomorphic structure and the apparent conventionality of pictures is inherent to the very logic of depiction – and as such does not merely concern propositional signs, but any state of affairs *qua* picture. In the framework of the T*ractatus*, even thoughts cannot escape their pictoriality. So to the question: what is

[11] Wittgenstein does claim that "a thought is a *proposition* with a sense" (*TLP* 4 my emphasis), but I think we should understand here "proposition" as "propositional sign". After all, a proposition is a propositional sign with a sense…

[12] On this point, there might be a useful parallel between the *Tractarian* contrast between symbol and sign – and the *Tractarian* contrast between thought and propositional sign…

common to all signs – i.e. pictures – that express the same sense? The only answer we are left with appears to be: that they are used as such. To be clear, I am not suggesting that the *Tractatus* somewhat already contains an explicit endorsement of Wittgenstein's later view that the meaning of a sign is its use (see for example *BB* 4-5 and *PI* §43). Rather, I am simply suggesting, once more, that such considerations are not altogether alien to the *Tractatus*. Quite the contrary.[13]

Within the framework of the *Tractatus,* there is nothing left but the use we make of signs to ground the fact that they express the same sense – and so that they represent or are a picture of the same state of affairs. But this is not a solution to the tension we have highlighted between the requirement of isomorphic structure and the conventionality of picture. It is rather a concession that such a requirement is, we could say, *fantastical*. Under our proposed reading, the *Tractatus* starts by fleshing out the picture theory and with it a central insight: relations of picturing are grounded on the formal – and so essential – properties of states of affairs. And the relevant properties of these states of affairs are their structure. But then, a concession is made: *qua* pictures, states of affairs do not have determinate – or fixed – structures. They do not have the relevant formal properties. So how can they still stand in relations of picturing? Well, if it is not on the ground of something essential – and so intrinsic – to them, it ultimately has to be something about how we used them…

[13] That is how, I think, we should understand the suggestion that a thought is a "propositional sign *applied* and *thought out*" (*TLP* 3.5 *my emphasis*).

If I am right then, we should read the *Tractatus* not as a sound edifice, that was only later found to be resting on a wrong foundation, but rather as an edifice already visibly cracked and on the verge of collapse.[14]

It is tempting to summarise the *Tractarian* conception of picturing as follows: the relation between a picture and what it represents is internal (i.e. the relation holds solely on the ground of the essential/intrinsic properties of both state of affairs). And the relation is not just an expression of how we conceive of the picture and of what it represents: it reveals (objective) features of both – in light of which they are isomorphic. Our proposal is that such a conception of internal relations holding *de re* – that is solely on the ground of the essential properties of the states of affairs themselves – does indeed find expression in the *Tractatus*. But, already in the *Tractatus* this conception proves to be largely fantastical.

3.3 Intentionality in *Philosophical Investigations*

There is often talk of a middle period between the *Tractatus* and *Philosophical Investigations*. And on the issue of intentionality, that middle period is often characterised as a progressive abandonment of the idea of a necessary isomorphism between a picture and what it represents, while still maintaining that the relation of picturing is *internal* (see Hacker 1996, Crane, 2010 and Blank 2007). It is also in this period that examples such as

[14] I cannot elaborate on the issue here, but my proposal also contrasts with so-called 'resolute' readings (see for example Conant 2002 and Diamond 1998 for early expositions of such readings). In keeping with the metaphor: I take it that the whole point of fleshing out the picture theory and its attending metaphysics in the *Tractatus* is not so we can attend to its collapse. That such a collapse is somewhat already announced in the *Tractatus*, I think there are much textual evidence for this. But that the collapse is planned from the start and serves an educational – or even liberatory – purpose? I find this less convincing. Why then spend so much time and energy on building the edifice, but so little in attending to its collapse, if the collapse is what matters?

that of an expectation and its fulfilment progressively make their first apparition and mark a shift of focus – or at least a new approach – to the question of intentionality. While I agree that work on this middle period can be enlightening, I think it also runs the risk of distorting our reading of *Philosophical Investigations*.

For example, the issue of the relation between an expectation and its fulfilment is still discussed in *Philosophical Investigations*. And even some remarks on it appear to have been taken almost verbatim from previous manuscripts dating back to this so-called middle period (on this point, see Hacker, 1996, 10-11). But it does not follow that Wittgenstein's position on the issue has remained unchanged from when these remarks were initially written, all the way to the composition of *Philosophical Investigations*. We should thus be careful in trying to shed light on *Philosophical Investigations* by appealing to manuscripts from the middle period. For this reason, I will focus mostly on *Philosophical Investigations* itself in what follows.[15]

In a passage that has sparked a lot of discussion, Wittgenstein wonders how one's expectation that someone will be coming can ever be fulfilled by that person actually coming. After all, as Wittgenstein insists, the expectation and its fulfilment, in this case, do not look anything alike. (To be sure, the sense of wonder would remain even if we were to conceive the expectation as essentially some kind of inner thought.). We thus face, once again, the issue that has been occupying us: in what way can these two things apparently

[15] A similar point could be made about attempting to shed light on the *Tractatus* by appealing either to the 1914-1916 notebook or to the manuscripts that were written afterwards. (This might explain, to some extent, the disagreement between my reading and Hacker's…). My position is that the *Tractatus* – just like *Philosophical Investigations* – should be seen as largely self-standing. The central claims of the *Tractatus* are expressed in the *Tractatus* – and that is where we should primarily be looking for them.

so dissimilar be internally related such that one can be said to satisfy the other? How can an expectation and its fulfilment be situated within the same space of possibility?

Wittgenstein ends up suggesting that "it is in language that an expectation and its fulfilment make contact" (*PI* §445). But then, what does this amount to exactly? It might be very tempting to read it as follows: one's expectation that someone, let's say N., will be coming over is fulfilled by N. coming over because the event of N. coming over, like the expectation in question can be expressed by the same proposition. There is a fit between one's expectation and the event that will fulfil it on the grounds of some kind of isomorphism – but an isomorphism that only manifests itself in language (i.e. in how we express the expectation and describe the event).[16] There are two ways to understand this suggestion. Under the first interpretation the expectation and its fulfilment fit together because they can be expressed – or at least thought of – in terms of N. coming over by some external observer. Under the second interpretation the event that fulfils the expectation does so precisely because the one expecting N. to be coming over is in a position to articulate their expectation and the event that fulfils it in these very terms. Of course, if the one expecting N. to be coming over can describe their expectation and the event that fulfils it in these terms, then so are the others with whom they share a language. This is why there is a sense in talking about how *we* would express these things. But this is precisely a case in which how one expresses their expectation and how *we* – who speak the same language – would describe it are aligned. The suggestion then is that an

[16] This is, roughly, Crane's reading (see Crane 2010 and 2014). Even though Hacker insists on contrasting his reading with that of Crane (see Hacker 2013a and 2013b), the contrast is largely irrelevant to our concern. After all, Hacker does say that, in *Philosophical Investigations*, Wittgenstein conceives of internal relations as holding only between concepts and propositions (see 1996, 31 and 32) and of language as autonomous (1996, 31 and 38).

expectation and its fulfilment fit together because, when having such an expectation or when ascribing it to others, *we* describe the expectation in question and the event that ultimately fulfils it in the very same terms.[17]

I think there are good reasons to be suspicious of the first interpretation according to which the fit between one's expectation and its fulfillment could be entirely independent of one's own linguistic capacities and be grounded instead solely on an external observer's distinct attitudes and expressions. One of the issues becomes quite obvious when we conceive the expectation that N. will be coming over as something like the thought that N. will be coming over – which I agree with McDowell, is neither problematic in itself nor obviously something Wittgenstein would deny (see McDowell 2013a). Indeed, if we acknowledge that the expectation that N. will be coming over is not just a series of actions (e.g. walking up and down the room, etc.), but also consists in entertaining the thought that N. will be coming over, then it is difficult to see how one could be having that very expectation without having the relevant linguistic – and conceptual – capacities to express the thought in those terms. The point is that it is not at all obvious how one could even entertain a thought with *that* content if one did not possess the linguistic – or conceptual – capacities to express their expectation as the expectation that N. will be coming over. And surely, it is because the expectation has *that* content that it is ultimately fulfilled by N. coming over – rather than by any other possible event.

[17] Now, I think this is very much in line with Crane's reading. Indeed, as he puts it, for the later Wittgenstein, "*we* find the "contact" between expectation and fulfillment in the fact that *we* use the same words ("he'll come in") as an expression of what *we* expect, and as a description of what fulfills it." (Crane, 2010, 95, my *emphasis*). Those expecting and those willing to express that expectation in terms of "he'll come in" are one and the same.

But even if my expectation does not consist in my entertaining the thought that *N. will be coming over* – that is, even if my expectation is not propositionally articulated –, another issue remains. If I am expecting N. to be coming over, I should be in a position to see N. coming over as the event that fulfils my expectation. It is not to say, of course, that I might not sometimes be confused as to what I really expect (e.g. I might be mistaken in thinking that my walking up and down the room or my being tense and my thinking about N. are manifestations of my expecting him to be over; they might rather be manifestations of my fear that his coming over will end up in an argument). The point is rather that, if I am indeed expecting N. to be coming over – and the content of my expectation is available to me –, then I should be in a position to appreciate the internal relation that holds between it and the event of N. coming over. Now, of course, if this internal relation is grounded on the fact that both my expectation and the event that fulfils it can be expressed by the same proposition, then it seems to follow that my expectation has to be available to me as having that propositional content. It cannot merely be describable in these terms by any external observer. So if the suggestion that an expectation and its fulfilment fit together because they are expressed – or described – by the same proposition is to even get off the ground, we must deny that the content of one's attitudes (e.g. of expectation) can be fixed independently of one's ability to express these attitudes

Properly understood, such a reading of an expectation and its fulfilment 'meeting in language' might not be self-defeating – and it might even be, at first sight, exegetically appealing. However, I want to suggest another reading which might ultimately be more faithful to the spirits of these remarks – and which, at the very least, might prove more fruitful given our current concern. As I see it, to show that our sense of wonder as to how

the expectation and its fulfilment can 'meet' is largely misguided in this situation, we can appeal to language's mediation in a rather different way. One worry with the preceding reading is that it appears to force upon us the thought that whether a particular event fulfils a particular expectation depends if not upon how we actually end up describing both, at least upon how we are inclined to judge and describe both. That is, it makes claims about the relation between the expectation and its fulfilment as we conceive of them rather than as they are in themselves. Even when conceiving the expectation as the thought that N. will be coming over, the proposal is that the relation of fit is grounded on the proposition we would use to express that very thought and to describe the event of N. coming over – not on the thought and the event themselves. According to such a reading, there is thus an important shift from the *Tractatus*: intentional relations (e.g. between a proposition and the state of affairs that makes it true, an expectation and its fulfilment, etc.), inasmuch as they are internal, cease to hold *de re* to only hold *de dicto*. But at the same time, such a reading appears to maintain the idea, central to the *Tractatus*, of a necessary isomorphism. Only, the isomorphism manifests itself, not in how things are, but in how we talk of them. What is lost then, if we accept such a reading, is the idea of intentional relations holding on the ground of something objectively there.

Of course, as we have discussed at length in the first chapter, such an idea can seem difficult to flesh out without falling into some kind of platonistic mythology (see for example McDowell, 1996 and 2000). But it is not impossible. And I think McDowell is right in suggesting that showing this is an important strand of Wittgenstein's later philosophy.

To appeal to any kind of isomorphism is, I think, ultimately to shy away from the issue that has been occupying us here. Our question has been: how can things apparently so dissimilar as a proposition and the state of affairs it represents, or an expectation and its fulfilment, ever be internally related? The move in the *Tractatus* is precisely to say that the dissimilarity is only *apparent*, that proper logical analysis will reveal shared essential features. Both the proposition and the state of affairs have, not only the same logical multiplicity, but also the same structure. But this is precisely to dismiss what motivated the question in the first place: on the face of it, the proposition and the state of affairs look nothing alike! And it is not at all obvious that Wittgenstein is dismissing this in *Philosophical Investigations* or trying to show it to be an illusion. As I see it, the illusion he is trying to dispel in these remarks is on the contrary the idea that we can only make sense of things being internally – and in this case intentionally – related if we can secure some sort of isomorphism between them. What Wittgenstein is attacking in this passage of *Philosophical Investigations* is precisely the idea that a proposition cannot represent a state of affairs if it does not have something in common with the state of affairs. Or the idea that an event cannot fulfil an expectation if we cannot uncover, below the appearances, something it shares with that expectation – whether this 'something' belongs to the expectation and the event themselves or only manifests itself 'in language'.[18]

[18] So, according to this reading, the dialogue that goes on in the second paragraph of *PI* §444 should be seen as *ironic*. After describing what N. coming over and expecting N. to come over might look like – and concluding that there is not the slightest similarity between the two events, the dialogue goes on: "So how can one use the same words in describing them [i.e. both events]? – But then perhaps I say, as I walk up and down: "I expect he'll come in." – Now there is a similarity here. But of what kind?!". I do not think Wittgenstein enjoins us here to be satisfied with the answer that I might utter the words "I expect he'll come" while indeed expecting N. to be come over. Of course, this would be a "similarity", something it has in common with the event of N. coming over. But, it misses the point: it is already accepted that I might use the same words to describe both events. The question is: how can I be justified in doing so? Noting that, in walking up and down the room, I might also utter the sentence "I expect he'll come" only pushes back the difficulty. After all, at first sight, the sentence "I expect he'll come" – or even the thought that he'll come –

Another worry might be the following: understanding an expectation and its fulfilment meeting in this way seems to force us to accept that non-linguistic animals might not have expectations at all. Surely a dog might expect its master to come home. Wittgenstein's point, as I take it, is not to deny that. It is not to say that expectations in general are only conceivable within the context of language. Rather it is that expectations, *for us*, are shaped by our initiation into a practice that is, in large parts, linguistic. That we are willing to describe one's walking up and down the room as one's expectation that N. will be coming over and the door opening and N. appearing in the doorway as N. coming over is not what ground the fact that the latter fulfil the former. Nor does the fact that we would, if queried, express our thought as the thought that N. will be coming over. It is not what makes it the case that the expectation and its fulfilment are intentionally – and so internally – related. Rather, that both can be described in terms of N. coming over reveals that such a relation holds between them. Had we not been initiated into the relevant practices, into the form of life that is ours, we might have never been into a position to *see* that my walking up and down the room is an *expression* of my expectation that N. will be coming over. We might have never been into a position to see N. opening the door and walking in as a satisfaction of that expectation. Ultimately, I might have never even been in a position to form such an expectation.

Now, of course, if we take Wittgenstein's quietism seriously, we should not expect to find in his later writings a substantive account of intentional relations. Rather, we should expect to find just enough to show that the possibility of such relations holding, in a variety of cases, is not really under threat. The sense of wonder as to how the event of N. coming

looks nothing like the event of N. coming over. And so, it raises again the question: "how can one use the same words in describing them?" (*PI* §444)?

over could ever fulfil my expectation is ultimately misguided. It does not call for our uncovering, below the appearances, some essential but hidden features. At the same time, it does not force us to accept that such a relation merely registers a correspondence of linguistic expressions.

What it forces us to recognise is rather this: if I am indeed expecting N. to be coming over, then I am in a position – i.e. I am appropriately equipped – to *see* N. opening the door and walking in as fulfilling my expectation. The mediating role of language in such a case operates well before I find myself expecting N. to be coming over. And I think, this is one of the central ideas of the later Wittgenstein: the practices to which we are initiated – and ultimately the form of life we take part in – are thoroughly and fundamentally shaped by language. Our linguistic capacities pervade our whole life to the point of informing even what is given to us in perception – or available in consciousness. (It is in that sense that it might be said that, for the later Wittgenstein, language shapes thought.) And this is *why* it can be said that "it is in language that an expectation and its fulfilment make contact" (*PI* §445). The fact that they so make contact is not merely a matter of linguistic conventions, but its availability to us is very much dependent upon our possessing the relevant linguistic capacities. And as we have already hinted, this dependence manifests itself in two ways. First, in making it even possible for me to expect N. to be coming over. And second, in disclosing certain events – like that of N. opening the door and walking in – as satisfying that expectation.

As I see it, such a reading fits quite nicely with some of Wittgenstein's remarks on the difficulty imagining non-human animals experiencing – or at least manifesting – certain attitudes we have, in contrast, no reluctance ascribing to human beings. As he puts it: "One

can imagine an animal angry, fearful, sad, joyful or startled. But hopeful? And why not?" And then: "A dog believes his master is at the door. But can he also believe that his master will come the day after tomorrow? – and *what* can he not do here? – How do I do it?" (*PI*, part II, section I, 1). As I understand it, what is denied here, is not the capacity of animals to experience – or manifest – intentional states *in general*. After all, beliefs surely are such states. And if there seems to be reluctance in attributing *some* beliefs to dogs, the reluctance is not about attributing beliefs as such. So, what is the difference between believing that his master to be at the door and believing that his master will come the day after tomorrow? Well, it seems to me a case could be made that the dog is appropriately equipped to enjoy the belief that his master is at the door – and ultimately to see the presence of his master at the door as confirming that belief. And if he is not in a position to enjoy the belief that his master will be coming the day after-tomorrow, it is because, further capacities are required to enjoy such a belief – capacities he ultimately lacks. In this latter case, the dog's limitation also concerns his ability to recognise the presence of his master at the door as confirming the belief he had that his master would be coming on that day.

An obvious difference in the two beliefs is of course the temporality. And it might well be that the dog is not appropriately equipped to believe – or entertain – anything so temporally distant.[19] But this limitation is, I think, secondary. What the dog crucially lacks, for Wittgenstein, is a language. And that is why he cannot believe anything that will be happening the day after tomorrow – or hope or expect for it to happen.[20] What about anger, fear or joy? Why is there no difficulty, for Wittgenstein, in imagining animals experiencing

[19] *PI* §650 seem to offer further evidence for the thought that Wittgenstein conceives dogs as generally lacking the relevant capacities to entertain intentional states towards temporally distant events.
[20] See the following passage: "Can only those hope who can talk? Only those who have mastered the use of a language." (*PI* part II, section I, 1).

or manifesting such attitudes? A plausible answer, I think, is that, for him, such attitudes *can* insert themselves in forms of life that are not shaped by language.[21] Hope, whatever its content, simply cannot.

It is worth noting that, when it comes to intentionality, Wittgenstein is not solely focused on intentional states (i.e. on mental states with intentional content). Much of what he says about the expectation of N. coming over and other such intentional states seems equally applicable to *meaningful overt signs*. For example, Wittgenstein raises the question of how an arrow can ever point (*PI* §454). I take the question as expressing the same kind of perplexity, we have already encountered: to wonder how the arrow can point in a direction is to wonder how the arrow might be internally related to that direction as a representation of it. And so, how the arrow and the direction might *fit* each other. (In the same way: to wonder how an expectation might somewhat anticipate its fulfilment, is to wonder how the latter fits – and so, is internally related to – the former.)[22]

Wittgenstein tentatively voices the following suggestion: it is because of its meaning that the arrow points in a direction. But he then promptly adds: "That is both true and false. The arrow points only in the application that a living creature makes of it". (*PI* §454). And again, there might be a temptation to read this as suggesting that the arrow points as a matter of mere convention – and so that the internal relation between the arrow and the direction it points holds solely *de dicto* on the ground of a kind of marching in order. If, for example, the arrow pointing north and the north direction were internally related, it would be because we agree on describing the former as "pointing north" and the

[21] This is of course not to say that a dog's fear of thunder and my fear of thunder are in any interesting sense 'the same'.
[22] On this, see *PI* §§437-439.

latter as "the north". Our proposal, in contrast, is that we can only *see* the arrow as pointing in the direction it points (e.g. the north direction), if we have been initiated into the relevant practices. But that the arrow is internally related to the direction in which it points is not dependent upon our taking the arrow to point in that direction, nor upon our treating it as such.[23]

So, in light of these considerations, there is another interpretation of Wittgenstein's later view on intentionality: in the case of an expectation and its fulfilment, or of an arrow and the direction it points, the relations are still conceived as internal – and claims about such relations as *de re*. The shift away from the *Tractatus* does not consist in a rejection of the insight that claims about such internal relations are *de re*. What Wittgenstein ends up moving away from is rather the thought that this could only be sustained by some form of isomorphism. In *Philosophical Investigations*, Wittgenstein is escaping the fantastical requirement at the heart of the picture theory – and the *Tractatus*. But he does so without altogether sacrificing the idea that internal relations involving intentional states are grounded on objective features of those states – and the world. And this view becomes available from the moment language – the use we make of it – ceases to be conceived as a set of more or less arbitrary conventions and becomes a world-disclosing achievement. To know a language in *Philosophical Investigations* is not merely to be able to navigate a network of grammatical rules and conventions. It is more fundamentally to be given access to and participation in a – particular – cultural world.

[23] This is a very schematic exposition. We will be coming back to this example in the next chapter (see infra 4.2.1 and 4.2.2).

Chapter 4:

Locating Bedrock; Laying the Groundwork

4.1 Introduction

We have suggested earlier that a satisfying account of rule-following would have to answer in some detail two questions, namely i) what features of signs are we *responsive* to in cases of rule-following? And ii) what is the nature of the relation between signs and our responses to them in cases of rule-following? Of course, not any answer would do. A satisfying account of rule-following would have to accommodate the insight that the features of signs we are responsive to are, in some sense, constitutive of signs themselves (we referred to this as the objectivity condition). It would also have to preserve the insight that the signs are, in some sense, guiding our responses to them (we referred to this, unimaginatively, as the guidance condition). Finally, and in contrast with McDowell's own proposal, we have suggested that such an account could be secured without already presupposing our responsiveness to linguistic – or at least conceptually structured – meanings. That is, we have suggested that such an account could be partially reductive.

We are now in a better position to outline what such an account might look like. And this is precisely what we will aim to do in this section. To be sure, we are still a long way from a fully-fledged account of rule-following: at the end of this section much work will still be needed. But our aim is somewhat more modest; we will have succeeded if we manage to sketch a cursory, but convincing answer to our two questions. Our aim is to lay the conceptual groundwork. It is to highlight how the objectivity and guidance conditions

can not only be reconciled - but reconciled in a way that is both philosophically substantive and partially reductive.

We have been coy so far as to what might justify our attempt to locate bedrock below the level of linguistic or conceptually structured meanings. Crucially, we have said very little about what might be *unsatisfying* about an anti-reductionist stance towards these meanings – and so why it might be important to depart from McDowell on the issue. In this section, we will try to offer some answers.

Finally, to conclude this section, we will suggest that our proposal is of interest not only because it can serve as the groundwork for a satisfying account of rule-following, but also because it implies a new way to think about – and ultimately defend – the so-called normativity of linguistic meaning and mental content.

4.2 What Features of Signs are we Responsive to?

1. Objectivity and Mere Conventionality

One of the central issues, so far, has been to circumscribe and ultimately secure an adequate notion of objectivity. In the preceding chapters we have argued, following Merleau-Ponty, that there are meanings that adhere to things in the world, that these meanings are not mere projections, and consequently that the corresponding intentional contents that are given to one in perception – or more generally in consciousness – do indeed track objective features of the world. If one perceives a tool as fitting a particular work, it is not because one attributes such a purpose to the tool – in light of one's practical intentions or beliefs and desires –, but because such a practical meaning is genuinely there to be perceived and, as such, is part of the fabric of reality. We have insisted that this

proposal is not only internally coherent, but that it also resists a variety of objections. We have repeated a similar line of thought in arguing that, for the later Wittgenstein, internal relations involving intentional states (i.e. states with intentional content) are indeed grounded on objective features of these states and their *relata*. One way in which we have expressed this is in saying that the internal relation that holds between an arrow and the direction it points – or between an expectation and the event that fulfils it – is not a matter of *mere conventions*. If language is involved in such cases, as Wittgenstein seems to believe, it is to the extent that its possession makes us sensitive to features of the world we would not have otherwise been sensitive to.

Still, in framing our proposal in terms of a contrast between being an objective matter and being merely a matter of convention, it might seem that we open ourselves up to a – potentially – damaging objection. However appealing our proposal, does it not end up denying what is precisely undeniable: namely, that the meaning of an arrow or the practical meaning of a tool is, in some sense, *up to us* or at the very least *dependent on us*? Surely, even though the arrow points in the direction it does, it could as well have been pointing in the opposite direction. And is this not precisely because there are no *intrinsic features of* the arrow that make it point in the direction it points – and correspondently no *intrinsic* features that would prevent it from pointing in another direction? Ultimately, if there are no such intrinsic features that constitute the arrow's meaning, then the meaning *must* be something extrinsic to the arrow: something we have come to agree on, have assigned to or projected onto the arrow.

Of course, we have repeatedly insisted, following McDowell and Wittgenstein, that our sensitivity and responsiveness to meanings is dependent upon our initiation into

relevant practices. But we have said very little about the genesis of these meanings. We have argued that they are there to be perceived, that they are part of the fabric of reality, but how do they come to be? If the answer is that we – or what we do – bring them into being, then is our talk of objectivity not empty air? And if the answer is that they come into being in complete isolation from us – we could say from our form of life –, then are we not back to denying the undeniable, namely our inescapable involvement?

Now, I think the objection is ultimately misguided – and we have already said a few words here and there as to why – but it is worth addressing it more explicitly and in some more detail. It highlights a widespread perplexity towards proposals such as ours; a worry that in trying to secure a demanding notion of objectivity here, we are either committing to a claim to objectivity we can never really substantiate or foolishly succumbing to a "magical" conception of signs' intentionality and meaning.[1] If our proposal is not to be dismissed outright, then we need to dispel this impression that what we are putting forward amounts either to an empty claim to objectivity or to a full-on "re-enchantment of nature".[2]

We need to show that there is a way to understand our involvement in the genesis of meanings that is compatible with them being intrinsic to – and so constitutive of – signs. The preceding rehearsal suggests that this cannot be done, that the recognition that meanings are, in some sense, dependent on us is *ipso facto* a concession that they are something we assign to signs rather than find in them. But the rehearsal is flawed.

[1] The phrase comes from Putnam (see Putnam, 1981, 3-5), but I think it voices a worry that is very much ubiquitous. Think of the insistence with which it is often claimed that the intentionality of overt signs simply *cannot* be intrinsic, that it *must* be derived from some other source, whether thoughts, communal agreement, etc. The most obvious – and influential - example might well be Searle in his *Intentionality* (see Searle, 1983, viii.)

[2] The expression, of course, is from McDowell (see for example 1996, 97).

What supposed fact about signs – and our responsiveness to them – justifies the thought that they *could* have meant something different? The rehearsal suggests something along these lines: if there is no difficulty in conceiving the same signs – the same physical marks, the same uttered sounds – as meaning something different, it is ultimately *because* signs are not the kind of things that have meanings *in themselves*. Now, this might seem obvious. But it is only obvious *assuming* a certain conception of signs – and more generally of objects. Of course, if signs are essentially dead, then something extrinsic to them *must* give them life (i.e. meaning). But why assume that they are essentially dead to begin with? The thought may be tempting, gripping even, but it certainly is not inescapable.

In discussing Merleau-Ponty, we have taken seriously his suggestion that we first and foremost inhabit a *culturally shaped* world – and so that what is given to us in perception and what we practically engage with ordinarily is a tapestry of objects bearing the mark of human action. We have focused on objects such as tools, but the same can be said of overt signs. Merleau-Ponty does concede that we can abstract from our culturally shaped world – and so, that we can "close our eyes" to the meanings that are made salient in it, to the point of treating the objects that surround us as mere bits of extension or of matter. But in doing so, we do not reveal what they *truly* are, in themselves. We rather forcibly "remove" or subtract one of their constitutive features (see Merleau-Ponty, 1945, 418).

Merleau-Ponty's point, then, is that we can conceive a tool – let us say a hammer – as merely a particular arrangement of metal and wood. And, so conceived, it will make no sense to think of the tool in itself as fitting or not a certain work. If it has any practical meaning, it will only be to the extent that we give it one – whether by convention or by

projection of our individual practical intentions. Obviously, the same particular arrangement of metal and wood could be used for any number of purposes. But a hammer is not just a particular arrangement of metal and wood. Just like an arrow is not merely an arrangement of lines.

I read Wittgenstein as making a similar point. In response to the suggestion that signs in themselves are dead, Wittgenstein offers the following analogy:

> It is as if we had imagined that the essential thing about a living human being was the outward form. Then we made a lump of wood into that form and were abashed to see the lifeless block, lacking any similarity to a living creature (*PI* §430)

Of course, if the 'outward form' of a sign (e.g. its particular arrangement of lines) is its only essential feature, then reproducing that 'outward form' will amount to reproducing that sign. The same goes for a living human being: if its only essential feature is its 'outward form' (i.e. what it looks like), then any lump of wood that shares such a form will *eo ipso* be a living human being. But, in that second case, we willingly acknowledge that it would be mad to think of the 'outward form' as the *only* essential feature of a living human being. A lifeless lump is not – and cannot be – a living human being, even if you somehow breathe life into it. Wittgenstein's point, I take it, is that it might be similarly mad to think of a sign as only an outward form, as a dead lump we need to force life – or in this case, meaning – into. Just like a living human body is essentially alive; signs may be essentially – and so constitutively – meaningful.[3]

[3] We do not have the space to elaborate on this, but it is worth noting that such analogies between the human body and (meaningful) signs come up quite often in Wittgenstein's later writings, as well as in Merleau-Ponty's *Phenomenology*. Merleau-Ponty insists on such analogies quite explicitly at times (see for example Merleau-Ponty, 1945, 197-198, 204). Wittgenstein is not as explicit, though I would contend that such analogies clearly transpire from his writings. (Take for example his response to the suggestion that we can never know if someone *really* is in pain, since we can never reach behind the pain-behaviour all the way to the pain itself: "Just try – in a real case – to doubt someone else's […] pain!" (*PI* §303). Compare with his response to the suggestion that to understand a sign, to grasp its meaning, we have to "get behind" the sign:

The difficulty here is to flesh out that insight without falling into a magical theory of meaning and intentionality. To the question "What gives [a sign] life?", Wittgenstein answers: "In use it *lives*", as if to make room for our involvement in the sign being alive. But, then, he goes on to ask "Is it there [in use] that [the sign] has living breath within it? – Or is the *use* its breath?" (*PI* §432).

Our involvement in the genesis of meanings – our role in signs coming to life – might be undeniable, but it remains an open question how best to flesh it out. Of course, if how we actually engage with a sign *is* its meaning, then our involvement in the genesis of that meaning will be incompatible with it being an objective matter. If the arrow's pointing amounts to nothing more than the fact that we actually use and respond to the arrow as such, then the thought that, in doing so, we are answerable to its meaning will become unintelligible. If our actual patterns of uses *determine* the arrow's meaning – if this is how we are involved in the arrow coming to life –, then there is little sense in saying that in following the arrow as we do, we are *responsive* to its meaning.

But there is another way to flesh out our involvement: instead of being straightforwardly meaning-determining, the actual use we make of a sign – and our response to it – might rather be conceived as the necessary context in which the sign can come to have the meaning it has.[4] The suggestion here is not so much that our actual use of the sign does not contribute at all to its meaning – it obviously does – but rather that it

"when I've asked someone something, and he gives me an answer (that is, a sign), I am content – that's what I expected – and I don't object: but that's a mere answer" (*PI* §503).)We can think of both the human body and meaningful signs as essentially *expressive* – and so as raising related philosophical difficulties.

[4] There might be a useful parallel here with Wittgenstein's insistence on the necessary role of "surroundings" in fixing the *content* of our representations. Wittgenstein gives the example of a coronation. It would not be a picture of pomp and dignity, if a crown was a parody of a respectable hat, gold a vulgar metal, the robes of the king, cheap to produce, etc. So, that our representation of a coronation is one of pomp and dignity – that our representation has that content –, is only possible in certain surroundings. Still, the content of that representation is not constituted or straightforwardly determined by these surroundings. (see *PI* §584).

does not do so in any constitutive way. There is no formula from which we can derive the meaning of a sign from its actual use (and vice versa). Now, if we understand our involvement in that way, we can still make sense of the meaning of a sign being "within it" – or, as we have put it: intrinsic to it. And so, we can still make sense of it being objectively there. There is, as it were, enough conceptual room between our engagement with a sign and its meaning.[5]

2. Creatures of Ours

In earlier chapters – and borrowing Brandom's phrase –, we have suggested that meanings are "creatures of ours". The phrase now reveals itself as being even more fitting. Obviously, if meanings are creature of ours, then it follows that, in some way, we *make them*. We are necessarily involved in their genesis. And so, this should certainly give reassurance that, in recognising the existence of meanings – and our sensitivity and responsiveness to them – we are not committing ourselves to a problematic re-enchantment of nature.[6] Meanings are not mysterious platonistic entities: they do not exist in complete

[5] Consequently, we can see how it is potentially misleading to reduce Wittgenstein's view to the slogan "meaning is use". Of course, such a slogan does seem to find expression in Wittgenstein's writings (I am thinking here especially of *BB* 4-5 and *PI* §43). And it seems undeniable that there is, for Wittgenstein, a tight link between the use we make of a word and its meaning. (We have conceded that much already.) But as *PI* §432 makes it clear: there is more than one way to substantiate that link… Another way to make essentially the same point: even if we were to accept the slogan at face value, there is still more than one way to understand "use". Let us say that the meaning of a sign is indeed its use in language. Should we understand "use" as *our actual patterns of uses* or rather as the function of the sign in our language? Of course, the issue we have raised appears a lot less pressing if we understand "use" in the second sense. The function of the arrow might be to indicate this direction, even if we fail to respond fittingly in all our actual engagements with the arrow. But ultimately this forces us into a similar position. The question then becomes: what is the relation between our actual patterns of use and the function of the sign? And for Wittgenstein, the answer appears to be: had our actual patterns of use been radically different, then the sign might not have the function it now has in our language. But the function of the sign is not just determined – or defined by – our actual patterns of use. So we are back at the point we were making…

[6] Or to put it differently: we are not succumbing to a kind of rule-as-rail mythology (see McDowell 1996 and 2000 and Wright 2001. See *supra* 1.4).

isolation from our engagement with them. And since we are effectively involved in their genesis, there is nothing queer or especially puzzling with them being within our perceptual and cognitive reach.

So, we have secured our involvement – or rather the intelligibility of our involvement – in the genesis of meanings. But in doing so we seem to be giving fuel again to the worry that meanings cannot be intrinsic features of signs – and so not objectively there. That is, if meanings are indeed creatures of ours and our involvement in their genesis is indeed pervasive, it seems to follow that, had we not been there – or had our form of life been radically different –, these meanings might not have come to be. Objects that are now meaningful (whether tools or signs, etc.) might have had different meanings or even no meaning at all. The source of the worry now is not so much that meanings might be too arbitrary, but that they might be too dependent upon the contingencies of our form of life. Still, the worry is essentially the same. As we have already conceded it appears undeniable that what we do is, in some way, involved in the genesis of meanings. But it appears equally undeniable that what we do, the kind of activities we engage in and how we engage in them is dependent upon our form of life. And so, any evolutionary or historical inflexion in our form of life could have led to us doing things differently – or even to us doing altogether different things. But then, if an object having the meaning it has is dependent upon what we do, then any inflexion in our form of life will potentially have repercussions on the object having indeed that meaning or not.

And so, we seem to face again the thought that any sign *could* have a different meaning than it now has – and that this possibility is ultimately tied to our involvement in the genesis of the meaning. Had we been initiated into a different practice or had our

physical constitution or natural reactions been in any way different, the arrow might have pointed in a different direction[7]. And this is true even if neither our actual practice, physical constitution or natural reactions are *constitutive* of the arrow pointing in the direction it now points. Ultimately, our involvement is such that, had we been different beings, there might have been different meanings.

On this point, there seems to be an important disanalogy with secondary qualities – one we did not acknowledge in our earlier discussion.[8] We did recognise that relating meanings to secondary qualities – in the way McDowell does – is useful in bringing forward the fact that certain features of reality might be inconceivable in complete abstraction from our sensitivity and responsiveness to them. And even though we did caution against taking the comparison too far, we did not push the matter, not quite having the conceptual resources to do so at the time. But we do now.

If we follow McDowell and think of secondary qualities as dispositions to elicit particular perceptual experiences in beings with the relevant perceptual apparatus (see McDowell 2011), then it is still conceivable that objects might have such qualities even in the absence of such beings. That is, a red chair might still be *red* (i.e. have the disposition to elicit the perceptual experience of *appearing red*…) even if, as a matter of fact, no being is relevantly equipped to perceive it as such. The chair might have such a disposition even if the disposition is never realised. So, the property of being red cannot be conceived in abstraction from its perceptual effect (i.e. its *appearing red*), but its possession by an object is not dependent upon its actually eliciting such an effect. And so, it makes sense to say

[7] Or to echo one of Wittgenstein's examples: had our natural history taken a different turn, it might have been fitting, to look at a pointing gesture from "fingertip to wrist", rather than the other way around (see *PI* §185).
[8] See *supra* 1.4

that the chair would still be red even if we had evolved in such a way as to be all colour-blind. It would still be red even if there were no beings – human or otherwise – to whom it would present itself as red. More generally, we could say that the quality of being red, the possibility for an object to possess that quality, is not dependent upon us – or any other beings' – having the perceptual, cognitive or practical powers we effectively have. Or again: it is possible for things to be red in complete isolation from any historical or – more likely in that case – evolutionary inflexions in our form of life.

But the same is not – and cannot be – true of meanings. We might not be involved in the chair being the colour it is – except maybe if we are the ones having painted it. But we are *necessarily* involved in the genesis of the arrow's meaning or the hammer's purpose. And so, the comparison with secondary qualities cannot, by itself, answer every concern we might have about the meaning of signs being an objective matter, about us finding meanings in the world. Meanings are creature of ours, but secondary qualities simply are not. And so, if the former are indeed part of the fabric of reality, it will not be in the exact same way as the latter.

As I see it, McDowell does not spell this out explicitly enough. He often insists that meanings are not constituted in complete isolation from us, but he rarely fleshes out the extent of our involvement. He often does say something about *how* we come to be sensitive or responsive to meanings, but he does not say as much about how meanings come to be. And that is why the distinction he makes between his own approach to meanings and rampant platonism is so often criticised as unsubstantial (see for example Wright, 2000). Claiming that meanings are part of the fabric of reality requires this further concession: bits of it are of our own making. And I think it is hopeless to try to avoid that concession –

and intellectually dishonest to try to diminish its import. What is required is a way to flesh out that concession in a way that does not deprive us of our pretention to objectivity.

My proposal here is that we need to take seriously the idea that meanings are *creatures* – rather let us say than *creations* – of ours. We are inescapably involved in their genesis, but as any creatures, they ultimately have a life of their own. The point of the metaphor is something like this: even if the meaning of a sign is dependent upon what we effectively do and ultimately on the trajectory of our natural history, it is still possible to make sense of us being mistaken about that meaning, of our failing to perceive it or respond to it in a fitting manner. We can be careless makers; we can fail to keep up with our creatures, etc. There is enough conceptual room for such possibilities. There is an inescapable link between the meaning of a sign and the kind of beings we are. And the link is tight enough to guarantee that such a meaning is, in principle, within our reach. But it is not tight to the point of rendering inconceivable our failure – whether momentary or repeated, individual or collective – to grasp it. The existence of a meaningful sign like the arrow depends on the existence of beings with perceptual, cognitive and practical powers somewhat like ours. But the possession of such powers does not, in itself, guarantee one's success in perceiving the meaning of the sign, nor one's success in responding to it fittingly.

What these considerations show is that the dependency of meanings on what we do – and ultimately on the kind of beings we are – still leaves enough room for their being objectively there. In conceding that, given a different trajectory in our natural history, certain meaningful signs might not have come to be, one is not *eo ipso* denying that signs are intrinsically meaningful. Obviously, this falls short of a demonstration that signs are indeed *intrinsically meaningful* or to come back to Wittgenstein's metaphor: essentially

alive. I am inclined to think that there is no such demonstration. At best, we can show that there is nothing especially mysterious in thinking of signs as essentially alive. That to think of them in this way is not to deny anything.[9] And certainly not our undeniable involvement in their being alive. This is not insignificant. Take the competing picture of signs as essentially dead. Our initial assessment was that its appeal lies primarily in the ease with which it accommodates our involvement in the genesis of meanings. It makes such matters unmysterious. But then, if we can show that the picture of signs as essentially alive can also do so with ease, then the competing picture of signs as essentially dead loses the edge it initially seemed to have. [10]

3. A Partial Reduction

Notwithstanding our claim that McDowell does not stress enough an important disanalogy between meanings and secondary qualities, there is a sense in which we are largely in agreement with him. Our position, just like his, implies the following: meanings are bedrock, such that phenomena involving meanings cannot be reconstructed out of materials that are not themselves already meaningful. Or to put the matter differently: living signs simply cannot be made out of essentially dead materials. Signs' meaningfulness, signs' being alive, is constitutively dependent upon there being such things as meanings. The difference is that, *contra* McDowell, we maintain that it is not essential to meanings that they be conceptually structured, whether in the sense of being of

[9] As I read them, this is a strategy Wittgenstein – and McDowell – often adopt.

[10] That McDowell, Merleau-Ponty or Wittgenstein may fail to provide a demonstration that signs are intrinsically meaningful is thus not really a failure at all. If there is a failure on their part, it is rather a failure to make clear how, in proposing such a picture of signs, one is not denying anything… The point here is not that there are such failures on their part – although we have, at times, expressed some dissatisfaction with certain aspects of their respective discussions. The point is rather that if there are such failures, we need to be clear on what they amount to.

the form of a concept, or in the more demanding sense of having been given linguistic expression (see *supra* 2.7).[11]

Unsurprisingly, the same goes for understanding: for us, as for McDowell, one's understanding of the meaning of a sign – or of an object – is bedrock. It cannot be made out of materials that do not already manifest understanding of – or sensitivity to – that meaning. But in contrast with McDowell, we maintain that the possession of such an understanding is not dependent upon concept possession. Our sensitivity to signs – and more generally to objects – in cases of rule-following need not be shaped or permeated by our conceptual capacities.

To be sure, the disagreement is not merely verbal: the difference between our position and McDowell's is not just that we mean different things by "meaning". Rather it is that we do not locate bedrock at the same level.[12] On that issue, as we have already discussed, we decidedly side with Merleau-Ponty.[13] But what should we make of Wittgenstein's own appeal to bedrock when it comes to rule-following? On which side of the disagreement might he be standing? Obviously, McDowell sees his own position as a faithful inheritance of Wittgenstein's (see for example McDowell 1984, 1992 and 2013b. See also supra 1.2). However, I think this can be disputed.

Take for example Wittgenstein's insistence that, when it comes to rule-following, the difficulty is "not, to dig down to the ground; no, it is to recognise the ground that lies

[11] McDowell would agree that meanings need not be conceptual in the second sense, but he would maintain that they are at least conceptual in the first sense (see McDowell, 2013c, 318-319).
[12] If one were to insist that meaning and understanding are essentially conceptual, in the spirit of McDowell, we could simply reframe our position as follows: when it comes to rule-following activities, we can indeed dig below the level of meaning and understanding, although we cannot dig 'all the way down'. (As a matter of fact, this is how we initially framed our proposal in chapter 1.)
[13] See *supra* chapter 2

before us as the ground." (*RFM* VI-31).[14] The point seems to be the following: we want to get clearer on what is going on when we engage in rule-following activities. But at the same time we are dissatisfied with how quickly we appear to reach the ground, the bedrock. So we convince ourselves that it must be possible to dig further down. Obviously, there is a quietist moral to the passage: it concludes with the suggestion that, on this issue – as, allegedly, on many others – "our disease is one of wanting to explain" (*RFM* VI-31). But it also seems to be making a positive philosophical claim. What grounds our rule-following activities, what our engaging in such activities consist in, does not require large excavation work to be uncovered, as it lies near the surface of the phenomenon.

Now, this is certainly compatible with McDowell's position: superficially, rule-following activities appear to involve our sensitivity to – or understanding of – meanings. And it turns out, for McDowell, that this is precisely what grounds such activities. There is nothing below these meanings and our sensitivity to them to be unearthed – or at least: there is nothing to be unearthed out of which we may hope to reconstruct them. But this is also perfectly compatible with our own proposal. In itself, Wittgenstein's appeal to bedrock, does not give us any clue as to where the bedrock effectively is, except that it is not as deep as we might be tempted to think. And this is true of our proposal as much as of McDowell's: in both cases we might be tempted to think that there is still room to 'dig deeper'.[15] But the temptation ultimately proves misguided.

So, while for McDowell, the fabric of reality has to be home to concepts – or at least to conceptually shaped things (see McDowell, 1996), it does not have to for us. But it has

[14] I think Wittgenstein is making a similar point in *PI* §217.
[15] There are, after all, other phenomena for which the bedrock proves to be far deeper – phenomena, which are not constitutively dependent upon meanings nor upon our sensitivity to meanings.

to be home to intrinsically meaningful things, to things which, in themselves, manifest a genuine form of intentionality. This is where our account proves to be *partially* reductive. As part of the fabric of reality there are things, meaningful things, with intentional content – and these cannot be reduced to things which are not contentful themselves. But meaningful things' contentfulness – or intentionality – is not dependent upon their being conceptually shaped – nor *a fortiori* upon their being given linguistic expression. *Contra* McDowell, our proposal thus raises the possibility that concepts could be built out of materials which do not presuppose concepts, but rather a more fundamental *form* of intentionality (see supra 2.7). It is reductive in that respect: it locates bedrock below the level at which anti-reductionist proposals, like McDowell's, would typically locate it. But the reduction is only partial: at no point is there an attempt to reconstruct intentionality out of non-intentional or non-contentful materials.

Our partially reductive proposal is thus a genuine alternative to McDowell's. But then, one might want to ask: why should we favour it? Why cannot we be content with the kind of anti-reductionism advocated by McDowell? Well, to begin with, it is worth noting that anti-reductionism about rule-following is often favoured, less for its own merit than for a perceived lack of any coherent alternative (see for example Boghossian, 1989, 548 or Miller, 2017, 23). Now, nothing in what we have discussed so far suggests that a partially reductive view of meaning and understanding – the suggestion of a genuine form of intentionality that is unmediated by concepts – is self-defeating or incoherent. And we have already shown that the most pressing objections to such a view can be, if not defused, at least answered. This is certainly not an argument for our proposal to be preferred. But it

shows at least this: there is an alternative to anti-reductionism about rule-following that needs to be taken seriously.

What might be an argument in favour of our proposal is the following: when it comes to our sensitivity and responsiveness to meanings, McDowell is forced to draw the line at competent human adults. Only those have the required conceptual capacities in his view. But this does not seem right. After all, are there not a lot of putative rule-following activities (e.g. playing games, classifying objects by colour or shape, using tools) that can be – and actually are – performed not only by children, but also by many non-human animals? McDowell is ultimately forced to deny that they have the relevant sensitivity and so cannot consider these cases as genuine cases of rule-following.[16]

The advantage of our proposal here is that it is agnostic as to where the line should be drawn. There is nothing in what we have said so far that implies that children – even young children – or certain non-human animals might not be sensitive to meanings and so manifest genuine understanding. That is, there is nothing, in our proposal, that prevents them, in principle, from being 'rule-followers'.[17]

---

[16] In the same way, given his view that perceptual experience is essentially conceptual, he cannot concede that children or non-human animals inhabit a world and has to say instead that they inhabit a mere environment (see McDowell 1996 and 2013c). To many, this is too high a price to pay (see for example Wright 2000).

[17] While I do not have the space to elaborate on the matter, I think the suggestion that certain non-human animals are rule-followers – and so are sensitive and ultimately responsive to meanings – is plausible in cases where they exhibit *culturally shaped* behaviours, whether such behaviours revolve around tool-use, social interactions or vocalization (for examples of such behaviour in non-human animals see Garland & McGregor, 2020 and Withen, Goodall & al., 1999).

When it comes to human children, it seems to me plausible that their sensibility and responsiveness to certain meanings might predate the acquisition of fully-fledged conceptual capacities. Take the following example: when encouraged to repeat "da-da", a child more often than not utters the sound "da-da". But at times they instead utter the sound "ma-ma" and then laugh. This latter response reveals, it seems to me, a kind of understanding of the respective (non-conceptually articulated) meaning of both "da-da" and "ma-ma". What makes the child's answer "ma-ma" funny – to them at least – is precisely that the meaning of "ma-ma" is somewhat polarly opposed to that of "da-da". Not any other sound would do for that purpose. We could say then that the child's answer is a joke, an attempt at humour. And this reveals understanding precisely because one can only playfully engage with the meaning of something, if one is first sensitive to that meaning.

Consequently, our proposal enables us to more easily maintain a *demanding* conception of conceptual and rational capacities. McDowell – and others – might well be right that only competent human adults possess such capacities. And our proposal makes it possible to defend such a view, without having to deny properly contentful (i.e. intentional) states to children and non-human animals. In contrast with McDowell, more traditionally naturalistic views of our conceptual and rational capacities often try to make the possession of such capacities less demanding – and so blur the line between, for example, discrimination of features of the environment and full-fledged concept possession. Part of the motivation seems to be, the perceived need to show the continuity between human cognitive capacities and that of non-human animals. But then, under our proposal, conceptual and rational capacities are parasitic on a more fundamental set of capacities (i.e. our sensitivity to and understanding of meanings). And if such capacities are shared with non-human animals – which, we have suggested, is perfectly possible – then the continuity is unproblematic even if we maintain a demanding conception of conceptual and rational capacities.[18]

[18] On this point, we might ultimately be departing from Merleau-Ponty. Obviously, the picture we are putting forward here of our (human) mode of engagement with the world is very much in line with Merleau-Ponty's own view on the matter. The point of departure rather concerns the status of non-human animals' mode of engagement with the world. Merleau-Ponty appears willing to recognise to them a sensitivity and responsiveness to some meanings, but he appears committed to a clear demarcation between the 'biologically basic' meanings perceptually given to non-human animals and the culturally shaped meanings that guide human action. This attitude is quite explicit in *The Structure of Behaviour*, but it also finds echoes in the *Phenomenology* (see for example Merleau-Ponty, 1942, 245-246 and 263 as well as Merleau-Ponty, 1945, 116). The insistence on such a clear demarcation evidently sits uneasily with the idea, we wish to put forward, of an unproblematic continuity between non-human animals and the human way of being in the world.
On this issue, I think Glendinning is right when he suggests that Merleau-Ponty ultimately endorses "the distinction, familiar in the work of Heidegger, Gadamer and McDowell, between animal inhabitation of an 'environment' or 'setting' (*Umwelt*) and 'man' living in a 'world' (*Welt*)" (Glendinning, 2007b, 112). What matters then is why. I would contend that Merleau-Ponty's endorsement is not *necessitated* by his view on the human mode of engagement with the world. It is rather brought about by his failure to appreciate the ways in which many animals may shape – and in turn inhabit – a cultural world. If I am not mistaken, this is also Glendinning's diagnosis (see especially Glendinning, 2007b, 115).

4. Kinds of Meanings

The feature of signs we are sensitive to in cases of rule-following is their meaning. Such meanings are "creature of ours" in an important sense – and in contrast with other potential parts of the fabric of reality, like secondary qualities. At the same time our involvement in the genesis of such meanings does not necessarily threaten their being intrinsic to signs. Drawing from our discussion of Merleau-Ponty we have further suggested that such meanings have intentional content, even though they are often not conceptually shaped nor propositionally articulated.[19]

Now, if our aim is the mere dissolution of a philosophical problem – or maybe perplexity – concerning the apparent tension between our involvement in the genesis of meanings and their being objectively there or concerning the required capacities for our participation in rule-following activities, we might be content to stop here. But if our aim is to offer a positive account of the phenomenon, then I think we will wish to say more. Obviously, not all directions of inquiry are equally promising here. For example, we have said very little about what is effectively going on in the genesis of meanings – and about what our involvement in that genesis ultimately amounts to. And so there might be a temptation for something more to be said here. But I am of the opinion that there is ultimately not much more philosophy can contribute on this particular matter.

When we suggested that our actual use of a sign might be necessary for it to come alive – become meaningful –, but added that our actual use was not *constitutive* of that life

---

Our suggestion then is that Merleau-Ponty's view on the human animal does not, in itself, speak in favour of his view on the non-human animals, and can even be used, *contra* Merleau-Ponty, to better appreciate the continuity between them.

[19] See *supra* chapter 2.

– and so not meaning-determining – we were making a grammatical claim.[20] Likewise, when we further suggested that there was no formula that gets you from the actual use of a sign to its meaning (and vice versa). The point we were trying to make was not that no such formula could be extrapolated from the *current* evidence – which would evidently raise the possibility that it might in light of *further* evidence. But rather that, since our actual patterns of use are not meaning-determining, no such formula is possible *in principle*.

Still, might we not want to know more about the phenomena by which signs becomes meaningful – and the role played by our actual use of signs in these phenomena? We might, but then what we are wishing for, I think, is not a philosophical account anymore, but an empirical one.

I think a more fruitful direction towards a philosophically substantive account of rule-following would be to reflect on whether or not there might be different *kinds* of meanings we are sensitive to and engage with. As we discussed earlier, Wittgenstein is very effective in his later writings at showing the great diversity of activities that falls under the umbrella of rule-following (see *supra* 1.6). But, given his quietism, he has little interest in attempting a finer-grained taxonomy of these different activities. Our proposal here is that such a taxonomy might be, not only possible, but philosophically informative. And it might revolve around such a reflection on what *kinds* of meanings we are sensitive – and ultimately responsive – to when we take part in these activities.

In what follows, I want to offer a brief sketch of what such a taxonomy might look like. It is certainly not exhaustive, but I think it goes over enough particular cases to show that, some generalisations are possible even if – and I think this is crucial – there is no

[20] On the idea of grammatical statements or claims, see *BB* 30 and 54-55.

unitary and entirely general account of rule-following. Ultimately, I think such a taxonomy perfectly exemplifies the unproblematic way of engaging in constructive philosophy we have sketched at the end of chapter 1 (see especially 1.7). It is not a fantastical attempt to uncover, once and for all, the *essence* of rule-following – and as such does not prejudge that there *must* be something common to all cases of rule-following. At the same time, such a taxonomy is more than just a *response* to specific philosophical concerns or perplexities (e.g. the regress of interpretations). It assumes, on the contrary, that philosophy can speak fruitfully on matters relating to rule-following, even if these matters are not already the source of any such concerns or perplexities.[21] Philosophy can take the lead, as it were, instead of being solely *reactive*.

5. Purposes

I want to suggest that in many cases of rule-following the feature of a sign – the kind of meaning – we are responsive to is something like its *purpose*. This is the kind of meaning we have already encountered in discussing tools. But I think it is also the kind of meanings one finds in signposts, orders, questions and a variety of other signs. Even in pieces in a game (e.g. the king in chess).

Let us take again the order "slab!" in the language game of the builder and the assistant. What is the assistant responsive to when, upon hearing the sound "slab!" he brings a slab to the builder? Well, to the meaning of "slab!", to the feature of the order that

[21] One may wonder: is the sort of taxonomy we are proposing really valuable because the intelligibility it offers is itself valuable? (This is how we have framed things at the end of chapter 1.) Or is it rather valuable to the extent that it preempts philosophical concerns and perplexities to which we may have otherwise succumbed? I am of the opinion that the contrast might not be as clear-cut as it seems, though I recognise that this warrants further reflection.

makes it a request for a slab – rather let's say than (i) a description of a slab or ii) a request for another sort of building bock.[22] This feature then, I want to say, is the *purpose* – or again, the *function* – of the order "slab!". In the same way as the hammer, in the workshop of the shoemaker, may *fit* or even *call for the shaping of a piece of leather on a wooden last*, the order "slab!", on the construction site of the builder and the assistant, may *fit* or even *call for the bringing of a slab*.

As we hinted earlier, an important aspect of the meaning of the hammer is that it has a practical meaning, a meaning that fits or calls for a particular action. I want to suggest that this may be an important aspect of *purposes*, more generally. It is quite obvious in the case of an order: ultimately, it does not call for something so much as for something *to be done*. The order "slab!" fits the action of bringing a slab, not just the slab as it sits there on the ground. I think the same goes with a question: it calls for an answer, for something to be said. The interrogative "slab?" while pointing at a slab, for example, calls for an expression – and so an act – of assent. A signpost, let us say a stop sign, calls *for* one to stop (i.e. for one to perform the action of stopping). The king in a game of chess fits a variety of moves, given the constraints that it cannot move more than one square and cannot put itself in check. In all such cases, there is thus an internal relation of fit that holds

[22] To be sure, in saying this, we are not committing ourselves to some version of the Frege-inspired force/content distinction (see Hanks, 2007, 141-142). The point we are trying to make is not that the meaning of the order "slab!" can be decomposed as a force (in this case the imperative force) attached to a representational content (in this case a representation of slabs). It is merely that the order calls for the *bringing of a slab*. There are thus at least two ways in which the assistant may fail to respond (fittingly) to the order: he may do something else than *bringing* a slab – if not, nothing at all. Or he might bring something else than *a slab*. But the distinction is not fundamental for us. Or rather: it is a distinction we may make in describing 'from the outside' the language-game of the builder and the assistant, but it is not a distinction that needs to be available to either the builder or the assistant. As I see it, the assistant may understand the order "slab!", but fail to understand the question "slab?" or the assertion "slab." And his failure need not be explainable as a failure to understand the interrogative or assertive forces. There need not be anything 'common to "slab!", "slab?" and "slab." that the assistant grasps in all three cases.

between an object – or sign – and an action, on the ground of the meaning of that object – or sign.

A second important – and I think related – aspect of purposes is their apparent dependency not only on what we do, but on what we do in the context of an *instituted* practice. The genesis of the meaning of the order "slab!" is tied, it seems, with the genesis of the language-game of the builder and the assistant *as a whole*. Likewise, the king's purpose – its point in the game – is doubly a creature of ours: it is both dependent upon how we actually move the piece in a game of chess and on our having instituted or created that game to begin with. Or to put the matter differently: that the king is the chess piece it is – and so, has the purpose it has – is contingent upon a particular context of use (i.e. upon our doing certain things). But part of that context is what we do with the chess pieces other than the king. That the king is the chess piece it is is partly dependent upon our engaging in the game of chess *as a whole*.[23]

Finally, a third aspect of purposes seems to be their lack of definiteness.[24] There is clearly something tentative, if not approximative, in expressing what the purpose of a sign – or object – is. The case of chess is, I think, rather informative here. Here is one popular way to conceive of chess as a rule-following activity: the game has an aim (i.e. to checkmate the adversary's king) and a series of definite rules constraining how one might reach that aim, including rules about how each piece can be moved on the chess board. It is a rule-following activity – or so the picture goes – to the extent that, in engaging in a

[23] We do not have the space to elaborate here, but such considerations might go a long way explaining Wittgenstein's suggestion that "it is only in a language that I can mean something by something" (see boxed remark following *PI* §35). Though whether or not he is right that this concerns our engagement with meaningful signs *in general* remains an open question.

[24] It is an aspect we have alluded to in discussing the practical meanings that are given to us in perception, according to Merleau-Ponty (see *supra* 2.3).

game of chess one is somewhat responsive to and ultimately guided by *these* rules. So, then in such a picture of chess-as-a-rule-following-activity, the best candidate for the king's or the rook's – or any other chess piece's – meaning is the rule that specifies their respective moving pattern.

The obvious difficulty here, as Dreyfus has repeatedly pointed out, is that with the exception of novice players, such rules are not really present in players' conscious engagement with the pieces on the board (see Dreyfus 2005 and 2006.). And it is in a way unsurprising: more often than not, chess players are not *merely* attempting to conform their moves to the rules of chess, they are trying to win the game! And the path to such a win is largely underdetermined by these rules. And so, this seems to challenge the thought that the meanings the chess players find adherent to the pieces on the board – and that he is ultimately responsive to – are just *these rules.*

The difficulty emerges, in part, from the following: the distinction between playing chess and playing chess *well* is as important as that between playing chess and merely conforming to the rules of chess. Or again, between playing chess and not playing chess at all. The picture of chess-as-a-rule-following-activity we have just sketched does not really have the resources to flesh out that distinction – nor *a fortiori* to account for its importance.[25] An alternative picture might be this: chess players are sensitive and ultimately responsive to meanings, *adherent* to the chess pieces. But these meanings are far richer than any specific rules of moving patterns as they are also permeated by the aim of the game and the pieces with each other. I think this is quite obvious in the case of the

[25] I take Cavell to be making essentially the same point when he says that "a certain mastery of the strategy of a game is as essential to being described as *playing the game* as mastery of its moves is." (Cavell, 1999, 306).

king: its purpose is indissociable not only from how it *can* move, but also from the fact, for example, that it can be put in check. But it holds for all pieces. They ultimately have distinct defensive or offensive purposes, depending, in part on their initial disposition on the board, on their affinities with other pieces' moving patterns, on their usefulness as decoys, etc. And so, not any move that conforms to the moving pattern of a piece *fits* that piece. When the chess players perceive the purpose of the piece aright, what is given to him is a relation of fit between the piece and *a good* move.

This seems to be an important feature of the rule-following activities just discussed: our success in performing them comes in degrees. And I want to suggest that this is partly due to the fact that the very meanings we are responsive to in performing these activities lack definiteness. Even if we are generally sensitive to such meanings, it is only ever partially. (They are not the kind of things we can grasp "at once".) At the same time, and more optimistically, our sensitivity to these meanings is also, in principle, perfectible. At times, the chess player improving his play may thus be explained in terms of his gaining a greater perceptual acuity of the meaning of the chess pieces – rather, let us say than in terms of his engaging in more and more intricate reasonings.[26] Likewise, the shoemaker becoming better at making shoes, might not just be a matter of him acquiring a greater dexterity: it might be a matter of him becoming more sensitive to the perceptual meaning of the hammer and of the other tools he uses. Or again: his greater dexterity in making shoes might not merely reflect a gain in motor control, but also a greater sensitivity to practical meanings.

[26] This would certainly fit the kind of phenomenological observations put forward by Dreyfus (see again Dreyfus 2005 and 2006).

If, as we have suggested earlier, the perception of such meanings is a form of understanding, then what we are, in effect, suggesting here is that the understanding of purposes (whether of the chess pieces, of tools, etc.) comes in degrees. To the extent that a rule-following activity is one *on the ground* of such an understanding, we can thus better appreciate why the distinction between performing an activity and performing it well might be so important in our thinking about rule-following.

In other cases, like shoemaking or the language-game of the builder and the assistant, the lack of definiteness does not manifest itself in difficulties of expression. For example, it seems easy enough to express the purpose of the builder's orders: they are simply requests for various kinds of building blocks. But, such expressions are, at the same time, uninformative. This becomes apparent when we think of how different bringing a building block might look like in different circumstances. In saying that the purpose of the order "slab!" is a request for a slab, one has said very little about what *it is* to bring a slab. Or again: one has not specified the definite list of particular actions that *fit* the order – or are called for by it. Nor has one given an indication as to what the hierarchy of fit is for these particular actions. For example, in turning three times on himself before walking back to the builder with the slab in hands, the assistant might indeed be performing the action of bringing a slab. And, so his action might *fit* the purpose of the builder's order. But there seems to be a clear sense in which, the action of walking straight back with the slab in hands would *better* fit that purpose (in most circumstances at least). Or to put it differently: it seems that the builder's order calls for the latter action *more* than the former.

The same goes with the purpose of the shoemaker's hammer: we have suggested that, in the shoemaker's workshop, it is to shape pieces of leather on a wooden last. But,

such an expression does not give us any clear criteria for what it *is* to shape a piece of leather on a wooden last – and so again: if it gives obvious expression to the purpose of the hammer, it gives us no definite list of the particular actions that fit that purpose. Nor does it give us any indication as to which action the hammer most calls for in the situation (e.g. which way of holding the hammer, of stretching the leather on the last, with which angle or strength to hit the leather, etc.).

The point might be obvious enough, but the question is what to conclude from it? Well, of course if the meanings we are sensitive – and ultimately responsive – to were entirely definite, then it might be tempting to conclude that there is a failure on our part here, that the expression we give to the purpose of the hammer or the order "slab!" is an unsatisfying *approximation.* But if we are right and such meanings are not definite, then the answer seems to be: perfectly clear criteria, definite list of actions, or undisputable hierarchies are never attainable *in principle*. Not because of any cognitive failure on our part, nor of any failure of expression, but because the meanings we are tracking are not the kind of things that can be so given, at once.

It is tempting to think of a meaning or concept lacking definiteness, in terms of it lacking sharp boundaries, which, in turn, forces upon us a number of metaphors (e.g. colour patches with discernible shapes, but blurred boundaries, an enclosure that is not fully closed, etc.)[27]. And so, if we go down that way, the idea of lacking definiteness ends up being reduced to a lack of definitional criteria, to the fact that, for some cases, it will not be clear whether or not they belong within the boundaries. What I want to suggest is that it is only part of the story. In saying that purposes lack definiteness, we are also registering

[27] Wittgenstein offers such examples (see for example *PI* §76).

the fact that, the space within the boundary is not perfectly homogenous: it has texture or depth. What lacks definiteness is not just the line drawing the boundary; it is also the space within that boundary. Actions that fit a sign – or object –, on the ground of the purpose of that sign – or object – can do so in various ways and to various degrees.

And so, to better perceive the purpose of a sign, to gain a greater understanding of it involves more than just better discriminating between actions that fit the sign, and actions that do not, in limit cases. It also involves better discerning (i) various ways in which an action might fit the sign – and so, satisfy its purpose – and (ii) the extent to which an action might fit the sign *better* than other similar actions.

6. Basic Patterns

I want to suggest purposes are not the only features – the only kind of meanings – we are responsive to. I think we need to also acknowledge that we are responsive to what I want to call *basic patterns*. I am thinking here of a variety of cases: the continuation of arithmetic series, or of decorative patterns, the classifying of objects by colour or shapes, the translation of written scripts into phonemes, etc.[28] Thus, for example, if the pattern "-- * -- * -- *" fits the continuation "-- * -- *", it is not because of the *purpose* of the initial pattern – of its function or its point in the kind of activities we engage in –, but because of a regularity. The pattern is genuinely *meaningful* in the sense we have sketched in our

[28] As we have seen, Wittgenstein discusses most, if not all these cases (see *supra* 1.6). Since then, they have been repeatedly invoked as paradigmatic examples of rule-following (see for example Ginsborg 2012, Kripke 1982, McDowell 2013b, Wright 2001.)

discussion of Merleau-Ponty: there exists a relation of fit between it and its continuation, but that relation of fit is not one of purpose in any obvious or interesting sense.

My proposal is that the distinction between purposes and basic patterns is not merely verbal. Nor is it solely justified by its usefulness in dissolving certain philosophical confusions or perplexities. The distinction is substantive: it tracks down genuine differences (i.e. differences that are grounded on the constitutive features of meanings). The most salient of these differences concern the three aspects we have just discussed: (i) their relation to action, (ii) their dependency on what we do and (iii) their definiteness.

First, whereas purposes ground relations of fit between a sign – or object – and an action, I want to suggest that basic patterns rather ground relations of fit between a sign and other signs – or more generally between an object and other objects. Take the decorative pattern "-- * -- * -- *". The relation of fit between the beginning of the pattern and its continuation is fundamentally one between a sign "-- * -- * -- *" and another sign "-- * -- *" – rather than between a sign and the *action of writing down* another sign. Similarly, if there is a relation of fit that holds between a heap of green beads and a further green bead, it is one that holds, first and foremost, between a heap of objects and a further object, rather than between a heap of objects and the *action of adding a further object*.[29]

Many authors on rule-following have been impressed by the insight that meanings are not, *in themselves*, action-guiding (see for example Glüer and Wikforss 2009 and Hattiangadi, 2006 and 2009). For them, it is obvious that our understanding of the meaning of a sign does not impose on us any kind of *obligation* or constraint as to how to act. Ultimately, the meaning of a sign does not even identify actions as fitting the sign or being

[29] The example is taken from Ginsborg, 2011b.

called for by it. Now, if we set aside for a moment the normative gloss of these considerations, I think we find a fair assessment, not of meanings in general, but of basic patterns. Such patterns are simply not action-guiding in the way purposes are. Their relation to action is not as straightforward.

We might still be able to make sense of a certain relation of fit between the beginning of a decorative pattern and the action of writing down its continuation, but that relation will be *derived* from the more fundamental relation of fit between the beginning of the series and the signs that constitute its continuation.

Here is one possible way to think about it: if the continuation "-- * -- *" fits the sign "-- * -- * -- *" on the ground of the basic pattern it instantiates, then it seems reasonable to think that it also fits the sign "-- * -- * -- * -- * -- *", and then the sign "-- * -- * -- * -- * -- * -- * -- *", etc. After all, all these signs instantiate the *same* basic pattern. And so, if one were to continue the initial sign by writing down, "-- * -- *", and then were to further develop the resulting sign by writing down "-- * -- *", etc., it seems one, in each case, would be doing the same thing, or be going on in the same way. Not merely in the sense that one would be writing down the same string of dashes and stars, but also in the sense that, in doing so, one would be *responsive* to the same basic pattern – and so to the same meaning.

In light of this, we might say that there is a relation of fit between the beginning of the decorative pattern *as it is developed* and the action of further developing it in the same way. Or again: that there is a relation of fit holding between the beginning of the *written down* pattern and the action of writing down its continuation. But this assumes that we can make sense of the initial string of dashes and stars as, itself an active development of the

basic pattern. Which, in turn, is only intelligible if there is already a relation of fit holding between the signs themselves.[30]

Second, unlike purposes, basic patterns do not appear to depend upon instituted practices. That the green bead *fits* with the other green beads is not dependent upon our engaging in and actively maintaining the practice of classifying beads by colour or any other such practices. So, even if we had never grouped objects by colour in that way, the relation of fit would still hold between the green bead and the other green beads.[31] The same is true, I would argue of decorative patterns – and even of arithmetic series. To be sure, the point is not that basic patterns are not dependent on us, on what we do or the form of life that is ours. It is rather, and more modestly, that contrary to purposes, such a dependency is not *also* a dependency on an instituted practice. (Think here of the way in which the purpose of the king appears to be dependent not only upon what we do with it, but also on the game of chess *as a whole*.)

The contrast between purposes and basic patterns here is not between meanings that are creatures of ours and meanings that are there independently of our involvement, nor is it between culturally shaped meanings and meanings faithfully tracking down natural kinds. It is simply an attempt to register two different ways in which meanings can be

[30] On this point, we decidedly side against Ginsborg. In our view, when it comes to decorative patterns or arithmetic series "going on" in the same way is not fundamental or "primitive", but very much dependent upon there being things with *intentional content* – in this case signs instantiating *basic patterns*. (For the most recent articulation of her view, which she also attributes to Wittgenstein, see Ginsborg 2020 and 2021 especially.)

[31] The point is rather straightforward, I think, if we imagine them as being all of the same hue, though I would concede that the matter becomes a bit more complicated if we imagine the beads as being all of slightly different hues. The meaning of the heap of beads of different hues of green might not be a 'pure' basic pattern. After all, whether or not a green bead of a slightly lighter hue 'fits' with the heap might depend in part on why we are grouping the beads in that way (and so, we could say, on the purpose of the heap).

creatures of ours or culturally shaped – that is, two different ways in which they can be dependent on us.

Thirdly, basic patterns are a lot more definite than purposes. There is a clear line of demarcation between, for example, the continuation that fits the beginning of a pattern and those that do not – or between the beads that fit the heap of green beads and those that do not. But there is more to the definiteness of basic patterns: whether or not a bead fits the heap of green beads or a continuation fits the beginning of a pattern is not a matter of degree: either such a relation of fit holds or it does not.

I think it is this aspect – the definiteness – of basic patterns that makes especially compelling certain images with which Wittgenstein grapples towards the end of the so-called rule-following considerations. I am thinking here of the image of grasping the meaning of a sign "at once", or "in an instant", of the meaning determining "the steps in advance" (*PI* §190), like "an ideally rigid machine" (*PI* §194). Or again, the image of the beginning of a series being like the "visible section of rails invisibly laid to infinity" (*PI* §218). The temptation to say: "All the steps are really already taken" (*PI* §219), etc. My suggestion is that all these images can be taken as attempts to capture the definiteness of basic patterns. And, when they are taken as such, they are quite inoffensive. It is a matter of not being then seduced into the kind of rules-as-rails mythology, we have previously discussed. It is also a matter of recognising that these images, if they can capture an important aspect of *certain* meanings (i.e. basic patterns), do not reveal something about all meanings.

These remarks on purposes and basic patterns are, to a large extent, still exploratory, though they hopefully make a compelling case for the fruitfulness of attempting a

taxonomy of meanings. They sketch a way forward, but in doing so, they inevitably raise a number of further questions – questions which, by themselves, certainly warrant further investigation. To begin with, if we are right that purposes and basic patterns are two distinct kinds of meanings, are there other kinds of meanings we are sensitive and ultimately responsive to? Furthermore, even if purposes and basic patterns are distinct, might one kind of meaning be more fundamental than the other? And if so, in what way? Finally, when it comes to linguistic signs, what kinds of meanings are we sensitive to in our engagement with them?

4.3 What is the relation between a sign and our response to it?

1. Motivation

Now that we are (a bit) clearer on (some) features of signs we are sensitive – and ultimately responsive – to in cases of rule-following, we can now address the question of the nature of the relations between signs and our responses to them in such cases. Our proposal is that they are relations of *motivation*. In light of our previous discussion of the notion – and of Merleau-Ponty's insights on perception and action more generally –, our proposal implies that, in cases of rule-following, an agent's interaction with a sign – or object – meets the following three conditions:

i. *the internal relation condition*: there is an internal relation of fit that holds between the sign and the agent's action

ii. *the causality condition*: the sign is causally involved in the bringing about of the action of the agent

iii. *the understanding condition*: The action is brought about on the ground of the agent's understanding of the relation of fit holding between the sign and the action.

One of the difficulties we encountered in discussing Merleau-Ponty's notion of motivation was to fully flesh out what it means for a motive and what it motivates to be *internally* related. We found it useful to reframe the problem in metaphysical terms: if such a relation indeed holds between a motive and what it motivates, what grounds it? What makes it so that the shoemaker's hammer, for example, motivates – and as such is internally related to – the realisation of a particular sort of work? We sketched the following answer: the hammer *fits* the work in question and it so fits it *because* of its meaning or intentional content. It is the hammer's meaning – as well as the nature of the work in question – that ultimately grounds the relation of fit.

The remaining unclarity concerned what makes such relations of fit properly *internal*. On this, we suggested, Wittgenstein proves to be a more fruitful interlocutor than Merleau-Ponty. Although Wittgenstein does not frame the issue in terms of motivation, he does reflect on the matter when he discusses intentionality – and internal relations involving signs with intentional content.

So, then the internal relation condition captures Merleau-Ponty's insight that a motive *fits* what it motivates; *calls for* it. And, in light of our engagement with Wittgenstein's discussion of intentionality, we can add that such relations are *internal* in the following sense: they are *solely* grounded on intrinsic features of both *relata* – in this case: the meaning, or intentional content, of the motive and the nature of what it motivates. It has the important implication that appreciating the kinds of action that fit a motivating sign – or object – is nothing over and above appreciating the meaning of that sign – or

object. Thus, for example, to appreciate the practical meaning of the hammer is *eo ipso* to appreciate the relation of fit that holds between it and a particular sort of work. To appreciate that relation is not an additional cognitive feat.[32]

The internal relation condition, if it is necessary, is certainly not sufficient: not all relations of fit underlie relations of motivations. A relation of fit might hold between a sign – or object – and an action without the sign motivating that action. Take an example we have already discussed: the expectation that *N. will be coming over* and N. actually coming over (see *supra* 3.3). Clearly, in this case, there is a relation of fit that holds between the expectation and its fulfilment. But, in most imagined scenarios of my expecting N. to be coming over, my expectation does not in any way contribute to the actualisation of N. coming over. Of course, there might be some scenarios in which that might be so – i.e. scenarios in which my expectation somehow brings about what I expect or at least, contributes to bringing it about. For example, if I let N. know that I expect him to be coming over, and if my letting him know is what convinces him to indeed come over, then it seems that my expecting him to come over has brought about his actually coming over.[33] But this is not a condition on expectations, in general. Take the alternative scenario in which I expect N. to be coming over, because I know my partner has invited him and N. has no idea that I know of the invitation. In this scenario, there is clearly no way my expecting N. to come over could contribute to his actually coming over. The point here is that the relation between an expectation and its fulfilment does not depend upon the former somewhat

[32] It would be an additional cognitive feat however if the relation between the motive and what it motivates was external: then one could be aware of all the intrinsic features of the *relata* and yet be unaware of the relation between them holding (see *supra* 3.2).

[33] Obviously, we are assuming here that I am letting N. know that I expect him to be coming over *because* I indeed expect him to be coming over. We are further assuming that, had I not let him know, he would not have come over – thus avoiding any potential issues of overdetermination. Under these assumptions, there is thus a clear sense in which my expectation sets in motion a series of actions that ends with N. coming over.

bringing about the latter. That this sometimes happens is incidental, and in no way constitutive of what makes a state one of expectation.[34]

But, as we have seen, such a constraint is crucial to Merleau-Ponty's characterisation of motives – and so, of relations of motivation. For him, a motive not only fits the action it motivates; it also "brings it about". In the relevant circumstances, a motivating sign – or object – "triggers" a response from us. It "moves us" to perform the action it calls for (see Merleau-Ponty 1945 see also Wrathall 2006 and 2007). In the workshop, the shoemaker's hammer is internally related to a particular sort of work, but it also has an effect on the shoemaker: it inclines him to perform that work. And, as we have argued, the most straightforward way to capture that insight is in terms of causal efficacy. To say that the shoemaker's hammer brings about the performance of a particular work by the shoemaker is to say that it causes – or at least is causally involved in – the performance. And this gives us the causality condition.

Taken together, the internal relation and causality conditions give us the following: a relation of motivation involves the overlapping of an internal relation of fit and of a relation of cause and effect between a sign – or object – and an action. As we have already discussed, this might prove to be a slight departure from the letter of Merleau-Ponty's characterisation of motives – and relations of motivations – but it is very much in the spirit of it. Obviously, *contra* Merleau-Ponty, this forces us to contrast relations of motivation with *merely* causal relations rather than with causal relations strictly. But this reframing

[34] This is not to say, of course, that my expectation might not bring about other kinds of action, for which a relation of fit might also hold. For example, my expectation might well *fit* a number of behavioral expressions of expectation (e.g. my walking up and down the room; my regularly looking at the clock, etc.). And it seems evident that, in most cases, my expectation would contribute to the actualisation of these behavioral expressions.

proves conceptually unproblematic and entirely innocuous to Merleau-Ponty's central insights (see *supra* 2.4).

The internal relation and causality conditions get us closer to the notion of motivation, but they still fall short of being sufficient conditions. Ultimately, it matters *how* the sign – or object – brings about (i.e. causes) the action of the agent. Tentatively, we can say that the causal role of a motivating sign must also minimally involve the agent perceiving the sign, being aware of it. Otherwise, there would be no sense in talking of the sign moving the agent to perform the action – ultimately no sense in even talking of the phenomenon as involving an agent proper. If a motivating sign triggers an action, it can only be to the extent that it somehow informs the performance of the action. And this is only possible if the sign is perceptually – and so epistemically – available to the agent. Thus, for example, in the workshop, the causal relation between the hammer and the shoemaker's work is grounded on the shoemaker's awareness of the hammer, to the extent that the awareness is informing his work. This contrasts with the brute ways in which a sign – or object – could trigger a series of bodily movements. A heavy box falling on the shoemaker might well cause him to fall. And there might well be a sense here in which the box *triggers* his falling, even though the relation between the box and his falling is not dependent upon his becoming aware of the box and his acting *on* that awareness. But then the box does not trigger or bring about the falling in the way a motive brings about what it motivates.

This is certainly part of the story, but it requires further refinements. After all, when it comes to relations of motivation, it seems to matter which feature of the sign – or object – informs the agent's performance of the action. That the hammer causes the shoemaker to

perform a particular work ultimately depends on his being aware, not just of any intrinsic feature of the hammer (e.g. its dimensions, weight, etc.), but also of its meaning. There might be a causal relation between the hammer and the shoemaker's work, even in the absence of such an awareness, but this would not underlie a relation of motivation. (Think here of a mechanical arm equipped with a number of visual and tactile censors for dimensions and weight, and which would perform the shoemaker's work solely on the ground of the perceptual availability of these features. The hammer might well be causally involved in bringing about the work here, but there would be no sense in saying that the hammer is *motivating* the work – even if we still assume the relation of fit holding between the hammer and the work.)

Thus, when a sign – or object – *motivates* the action of an agent – rather than merely and disjointedly fitting and causing it – the agent's r*esponsiveness* to the sign is grounded on the agent's *sensibility* to the meaning of that sign. Obviously here, we move quite freely from the agent being informed by the meaning of the sign, to his being sensitive to that meaning (and vice-versa). But, if, following Merleau-Ponty, we are right that such a meaning can be given in perception – and that the grasping or understanding of that meaning does not have to consist in anything more than that perceptual givenness - such a move is warranted. To perceive the meaning of the sign – or object – is, in itself, a form of understanding and so need not involve any further cognitive step in order to inform an agent's action. (see *supra* 1.4).

And if, as we have suggested, appreciating – and so perceiving – the meaning of a sign is *eo ipso* to appreciate the relations of fit that hold between the sign and various actions, these considerations give us the understanding condition. For a relation of

motivation to hold between a sign and an agent's response to it, it is not sufficient that the sign both fits and causes the response to it. It is also necessary that the causal involvement of the sign be dependent upon the agent's appreciation or understanding of the relation of fit.

2. Guidance

As we have repeatedly argued, a satisfying account of rule-following must meet the guidance condition: it must substantiate the insight that, in cases of rule-following *proper* – and in contrast with cases of *mere conformity* to a rule (or purpose or pattern, etc.) – the agent is somehow guided in acting the way he does. Our proposal then is that the notion of motivation offers us a way to substantiate that insight.

Let us come back, once more, to the language-game of the builder and the assistant. Our proposal unfolds as follows: the assistant bringing a slab in response to the builder's order "slab!" is a case of rule-following if and only if his action of bringing the slab is motivated by the builder's order. And such a relation of motivation implies that:

i. The builder's order "slab!" *fits* or *calls for* the assistant's action of bringing a slab
ii. The builder's order is causally involved in bringing about or triggering the assistant's action (i.e. his bringing of a slab)
iii. The assistant performs the action (of bringing a slab) on the ground of his understanding of the relation of fit that holds between the builder's order and that action.

As we have previously argued, the relation of fit that holds between the builder's order and the assistant's response to it is grounded on the *purpose* of the order. And the understanding

that grounds the assistant's response is ultimately an understanding of that purpose. The last remaining feature of our proposal is that this understanding need not involve any further cognitive steps in addition to the assistant's perception of that purpose.

If any of the three conditions is not met, then the assistant's action in response to the builder's order does not count as a case of rule-following.[35] But the third condition is really what distinguishes cases of rule-following *proper* from cases of mere conformity to a rule (or purpose, or pattern, etc.). If the assistant's response (of bringing a slab) were to be merely caused by the builder's order, it would still conform to – or be in accord with – the purpose of the order. That is, it would still fit or be called for by that order. But the assistant's response would not be *guided*, in any interesting sense, by the builder's order. If the third condition was not met, there would be no substantial difference between the assistant's response and the "jack-in-the-box" response of, for example, an automaton.

Now, how might our proposal go when it comes to actions involving our sensitivity to basic patterns, rather than purposes? It all depends whether or not a relation of fit between a sign and an action can be somehow derived from the sign instantiating a basic pattern. As we have argued, there might be a straightforward relation of fit between, for example, the beginning of a decorative pattern and a particular continuation of it. But this

[35] We have often talked not only of *cases* of rule-following, but also of rule-following *activities*. If, more often than not, the builder and the assistant's interactions (i.e. the orders and the responses to them) did not meet all three conditions, then, I want to say, the pattern of such interactions would not count as a rule-following activity. The idea here is that, an activity only counts as a rule-following activity if, more often than not, *our* participation in that activity is – or would plausibly be – underpinned by relations of motivation. Playing the language-game of the builder and the assistant is a rule-following activity – or counts as one – because, were we to play the game as the assistant, our responses to the builder's orders would plausibly be *motivated* by those very orders. The criterion we are proposing might appear rather weak: after all, according to it, a particular pattern of responses (e.g. regularly bringing a slab in response to the order "slab!") can be part of a rule-following activity even if it is not *necessary* for the actualisation of these responses that they be motivated. But the concession is crucial if we are to accommodate the insight that an automaton could be such that its pattern of responses to the builder's orders would be both merely caused by these orders and (behaviorally) undistinguishable from the pattern of responses of the assistant.

amounts to a relation of fit between a sign (e.g. ""-- * -- * -- *") and a further sign (e.g. "-- * -- *"). The question is whether or not there is also a relation of fit between the initial sign and the action of producing or writing down its continuation. We have offered some tentative arguments in favour of such a view. If they sustain closer scrutiny, then, there will not be any other significant difficulty in substantiating the insight that actions involving our sensitivity to basic patterns are underpinned by relations of motivation.

So, if there is indeed a way to derive a relation of fit between a sign and an action on the ground of a *basic pattern*, then our proposal enables us to account for most paradigmatic cases of rule-following: from an agent's response to an order, to his move in a game, all the way to his continuation of a mathematical series, or of a decorative pattern. But it does more than that: it also affords us the opportunity to expand the list of performances that are best classified as cases of rule-following. To begin with, if Merleau-Ponty is right in his characterisation of our use of tools – as we have argued he is –, then there is a strong case for their inclusion on the list. If the relation between the hammer and the shoemaker's use of it is one of motivation, then why would it not be a case of rule-following?[36] The argument for inclusion here is strengthened by the fact that the meaning of the hammer is of the same kind as the meaning we find in paradigmatic cases of rule-following like the assistant's response to the builder's order. In his workshop, the shoemaker is sensitive – and ultimately responsive - to the *purpose* of the hammer in the same way the assistant is sensitive – and ultimately responsive - to the *purpose* of the builder's order. There are other performances that Merleau-Ponty would unhesitatingly

[36] Relatedly, if the shaping of a piece of leather on a wooden last is *more often* than not underpinned by the hammer motivating us to perform that work– or at least those of us who have been relevantly trained –, then why would it not count, more generally, as a rule-following activity?

characterise as motivated, which, in contrast, we might be initially reluctant in classifying as cases of rule-following. I am thinking here, for example, of the suggestion that a pavement – inasmuch as it bears the mark of human action and so is, in itself, meaningful – could motivate a particular gait, or the suggestion that a particular political or social situation could motivate a revolutionary response (see *supra* 2.5-2.6). But I think it would be mistaken to give too much importance to our initial reluctance here. What matters is whether or not, in such cases, our response is grounded on our sensitivity to a genuine meaning – and so, whether or not, we are *guided* in our response by the pavement or the political or social situation.

There might seem to be a tension between the generality of our proposal here and our earlier resolve to stay clear of any ill-advised attempt to uncover, once and for all, the essence of rule-following (see *supra* 1.7). By proposing that a performance or action is a case of rule-following if and only if its realisation is underpinned by a relation of motivation, are we not, in effect, offering a maximally general account of what it is to follow a rule? And so, are we not succumbing to what we previously characterised as a misguided conception of what philosophy can achieve?

Ultimately, I think the tension is only apparent: in its details, our proposal proves coherent with our previous methodological commitment – and our criticism of traditional philosophy. First, despite its generality, our proposal makes it clear that relations of fit – and so, relations of motivation – can be grounded on many different *kinds* of meaning. We have discussed *purposes* and *basic patterns*, but we are very much open to the possibility that there might be other *kinds* of meaning that ground relations of fit and to which we are sensitive and responsive.

Second, and relatedly, we do not pretend that, in its current state, our proposal amounts to a *definitive* account of rule-following. We are adamant that it may lay the groundwork for a *substantive* account of rule-following, but such an account does not need to be definitive, settled once and for all – nor does it need to aim to be such. In attempting a taxonomy of meanings and in reflecting on what grounds relations of motivation, we are offering a number of tentative generalisations, but at no point are we erasing the diversity of cases of rule-following, the many ways in which a relation of motivation might hold between a sign – or object – and an action.

3. Normativity

At the beginning of our investigation, we were happy with McDowell's characterisation of the regress of interpretation as emerging from a conception of signs as being, in themselves, "normatively inert" (see *supra* 1.2). However, we have become progressively more wary of framing the regress of interpretation – or any issue relating to rule-following – in explicitly normative terms.

We are not disowning our previous characterisation, but we can now refine it in two important ways. First, and in contrast with what our previous characterisation might suggest, we can now say that it is not necessary for the meaning of a sign to impose any *normative* constraint on our responses to it. Whether or not our responses to a sign qualify as cases of rule-following does not hinge on that. What matters is that our responses are grounded on our understanding the meaning of the sign – and so, on our understanding the relation of fit that holds between the sign and various actions.

There is a clear sense in which an action that fits a sign is a "correct" action, one that "accords" with the sign. Hence, if McDowell is right – as I think he is – that ideas of correctness and accord are normative ideas, then it seems that a relation of fit would bear witness to a normative constraint. But even if it was not so: even if ideas of correctness and accord were not normative in any interesting sense – as many have recently argued (see for example Boghossian 2003 and 2005, Glüer and Wikforss 2009 and 2015, Hattiangadi 2006 and 2009) – our proposal would not be put in jeopardy. That is: even if relations of fit did not reveal any genuine normative facts and only had the appearance of being normatively constraining, understanding of such relations would still be sufficient to ground relations of motivations – and ultimately cases of rule-following.

Following Kripke's discussion of Wittgenstein's regress of interpretation, a lot of attention was given to the idea that cases of rule-following have to be dependent upon our grasping of some normative fact. As Kripke himself puts it: meanings are such that they determine "a unique thing that I *should* do" (Kripke, 1984, 24). But I think this somewhat misses the point. The focus should be on the idea that meanings can *tell us* s*omething* – and so guide use –, rather than on the idea that obligations can be derived from them.

This leads us to the second important refinement: we will be saying some more about this in the next section, but if relations of fit are indeed *normative*, I want to suggest that their normativity is *sui generis* and cannot be derived from other forms of normativity. If relations of fit do indeed give rise to normative constraints, such constraints will not be reasons to act, obligations or standards of what is good. Nor will they be derived from reasons, obligations or standards of what is good. They will participate in a distinctive kind of normativity.

Chapter 5

Following a Rule: The Contemporary Debate

5.1 Introduction

In *Wittgenstein on Rules and Private Language*, Kripke famously takes a stand against dispositional accounts of meaning. He invokes various difficulties such accounts have in answering the regress of interpretations *as he understands it*.[1] But for Kripke, the most decisive objection comes from the fact – which he takes to be uncontentious – that the relation between meaning and use is normative and cannot, as such, be captured by dispositions. That is to say, for Kripke, any attempt to reduce the meaningful employment of a sign to a regularity in behaviour – that is, to an actual pattern of use – is not merely open to challenges, but ultimately misguided since it overlooks an essential feature of the phenomenon (i.e. its normative character). Arguably, any other reductively naturalistic account of meaning would face similar limitations and thus would show itself to be equally misguided. Given such a radical implication, it is not surprising that the normativity thesis has been so widely discussed. Commentators were, at first, largely sympathetic towards Kripke's objection – and the insight on which it hinges –, many of them seeing the normativity of meaning as self-evident.[2] But it is a different story today. The normativity

[1] It does not matter for our current concern whether or not the paradox about the possibility of meaning Kripke attributes to Wittgenstein is indeed Wittgenstein's. Or whether or not the way in which Kripke introduces the paradox is faithful to the spirit in which Wittgenstein introduces the regress of interpretations. Consequently, we will not take a stand on the question.

[2] See for example, Baker & Hacker 1984, Blackburn 1984, Boghossian 1989, McDowell 1984, McGinn 1984, Wright 1984.

thesis has grown increasingly suspect and many arguments have been put forward in recent years to show either that it is false or that it does not have the consequences Kripke thought it had for dispositional – and other reductively naturalistic – accounts of meaning.[3] While some have suggested that, while false in the case of overt signs' meaning, the thesis might still be plausible in the case of mental states' content (see most notably Boghossian 2003), even that suggestion has been seriously challenged. As it stands, it cannot simply be assumed anymore that the meaningful employment of a sign – or the enjoyment of a contentful mental state – is normative in a way that threatens its naturalisation.

Our aim in this chapter is twofold. First, it is to situate the proposal we made at the end of the previous chapter within the context of this debate (see *supra* 4.3.3). Second, it is to compare and contrast our proposal with the original account of meaning and rule-following recently put forward by Ginsborg in a series of papers. (see Ginsborg 2010, 2011a, 2011b, 2012, 2020 and 2021). Ginsborg is confident that the notion of primitive normativity – around which her account of meaning and rule-following revolves – can not only ground the normativity thesis, but also answer Kripke's take on the regress of interpretations. In what follows, we will focus mostly on the former pretention of Ginsborg, though we will also say a few words about the latter.

There is an obvious rejoinder between Ginsborg's account and our own proposal: both argue for the intelligibility of cases of rule-following independent of the possession of fully-fledged conceptual capacity – and so, independent of the grasp of propositionally articulated 'rules'. But, while Ginsborg's notion of primitive normativity helps her in

[3] Recently, Boghossian and the Stockholm School have been among the most prominent figures of anti-normativism (see Boghossian 2003 and 2005a, Glüer and Pagin 1999, Glüer and Wikforss 2009 and 2015, Hattiangadi 2006 and 2007, Wikforss 2001).

making that argument, it is also the source of a number of difficulties – difficulties our proposal simply does not face. In light of this, we will argue that her account, inasmuch as it revolves around the notion of primitive normativity, is not the most suitable to substantiate the normativity thesis.

In a few passages, Ginsborg hints at a neglected, but in our opinion important, connection between understanding and intelligent action (see for example Ginsborg, 2012, 128). To conclude, we will say a few words about how our proposal might enable us to flesh out that connection.

5.2 Anti-normativism: a Brief Overview

According to the anti-normativists, there have been two main strategies to substantiate the normativity thesis. The first strategy has been to try to directly derive norms from the meaning of overt signs or the content of mental states. It goes as follows: the fact that an agent's utterances are endowed with determinate meanings – like the fact that the agent's thoughts have determinate contents – imply, it seems, that they are answerable to specifiable norms. For example, once it is established, that an agent means *green* by "green" – or that his thinking "green" has the content *green* – then his patterns of thought and behaviour involving "green" can be judged as correct or incorrect, or as fulfilling or violating specific obligations. And such judgments, it seems, boils down to whether or not the object referred to – or thought of – is indeed green. Whether we think of semantic criteria of correctness or of obligations governing assertions or beliefs, the norms posited ultimately appear to discriminate between the applications of "green" that are true and those that are not. As such, the judgement that an agent has misapplied a

linguistic expression or that he has failed to believe what he ought to can be made whether or not the agent recognised the norms his utterances and thoughts are submitted to. His believing, for example, that a red fire hydrant is green can be assessed as contravening the obligation of believing only what is true, whether or not he is aware of this obligation – and *a fortiori* whether or not he commits himself to abide to it. Now is this sufficient to ground the normativity thesis? It would be, or so the story goes, if such norms were both to express a genuinely normative evaluation and be derivable solely from signs' meaning or mental states' content. But, as the anti-normativists forcefully insist, none of these 'norms' satisfy both constraints. For example, there might be semantic criteria of correctness, but, understood in terms of truth, they are not really 'norms': they are merely the expression of an act of classification (between things that are green and things that are not, etc.).[4] As such they can be fully accounted for in non-normative terms. In contrast, agents might incur an obligation to speak only correctly. However, this obligation is pragmatic and prudential rather than semantic.[5] After all, it seems that a speaker can incur an obligation to lie; to apply, for example, "green" to something that is not, and still mean what he usually does by 'green'. If so, the obligation, for a speaker, to apply "green" only to green things does not follow from the fact that he means *green* by "green". For many, similar considerations speak against the suggestion that obligations, such as the obligation to only think "this is green" of things that are indeed green, can be derived from the specific content of one's thoughts.[6]

[4] On this point see Bilgrami 1993, Glüer and Wikforss 2009 and 2015, Hattiangadi 2007 and 2009, Wikforss 2001.
[5] On this point see Glüer and Wikforss 2009 and Hattiangadi 2007 and 2009.
[6] To circumvent this objection, Boghossian suggests that the normativity of mental content could be derived from the normativity of beliefs. He concedes that we might not incur an obligation to think only what is true. After all, we can imagine or desire things that are not the case (i.e. that the red fire-hydrant is green, etc.). But surely, Boghossian adds, we incur an obligation to only *believe* what is true. We could thus salvage the

Now, I think we can agree with the anti-normativists' conclusion here. There is indeed an issue with appealing to obligations (on beliefs or assertions) to substantiate the normativity thesis: it necessarily leaves a gap between the agent's appreciation of the meaning of their utterances – or the content of their thoughts – and the agent's appreciation of the norms their utterances or thoughts are submitted to. And this is seemingly irreconcilable with the idea that it is the very meaning of the agent's utterance – or the very content of their thought – that imposes the normative constraints. After all, if that was the case, how could the agent *really* appreciate the meaning of their utterances and not appreciate the normative constraints it imposes?

As for appealing to conditions of correctness, there is indeed an issue *if* we think of them as primarily – if not exclusively – tracking down truth. Then, as the anti-normativists point out, it would be quite possible to express them in non-normative terms. The normative gloss of conditions of correctness would only be apparent. But is this the right way to conceive conditions of correctness?

Our proposal is that it is not: conditions of correctness do not primarily track down truth, but rather *sui generis* relations of fit. And since, under our proposal, relations of fit are grounded on the meaning of signs – whether overt signs or signs that are 'thought out' – then there is no gap between appreciating – or understanding – the meaning of a sign and appreciating the relations of fit the sign is involved in. To appreciate the former is *eo ipso* to appreciate the latter. If relations of fit genuinely reveal normative constraints, as we have suggested, then we are some way towards a defence of the normativity thesis.

---

normativity thesis in the case of mental content by showing that thoughts in general (and so any kind of propositional attitudes) are conceptually dependent upon beliefs. (On this point, see Boghossian 2003).

Now, there is an obvious rejoinder between our proposal and the second strategy to substantiate the normativity thesis, criticised by the anti-normativists. But there is also a crucial difference. Indeed, according to the anti-normativists, the second strategy has been to argue that norms *determine* meaning and content – rather than being *derivable* from meaning and content. It revolves around the insight that the meaning of one's utterances – like the content of one's thoughts – depends on certain norms one has internalised and ultimately committed to. Such norms are like rules that guide the agent and in light of which they engage in various patterns of use. What the agent means by "green" – or what content "green" has in their thinking – is determined by the rule they are following in uttering – or thinking – "green". The agent could not employ linguistic expressions meaningfully – or enjoy contentful mental states – without recognising, from their own perspective, that they are bound by these rules. Consequently, the norms (i.e. the rules) that guide the agent's utterances and thoughts can ultimately serve as justification for them.

Postulating such norms circumvent the objections previously discussed and, as such, might appear to offer a promising way to substantiate the normativity thesis. Rules so conceived both determine the meaning of the agent's utterances and offer justification for these utterances. And so, just like our proposal, the second strategy appears to bridge the gap between signs' meanings and the normative constraints these signs are submitted to. That is, of course, assuming that such relations of justification are indeed normative.

However, and as we will see in the next section, there remains a crucial difference between our proposal and this second strategy. In our view, it is the meaning of a sign that grounds – or determines – the relations of fit the sign is involved in – and so, the normative constraints it is submitted to. The second strategy has it the other way around: it is the rule

expressing the normative constraints which grounds – or determines – the meaning of the sign. And this has fatal implications: given some widely held assumption about the nature of such rules, the second strategy ultimately falls prey to various versions of the regress of interpretations.

5.3 The Intentional View: Rules and Regresses

The second strategy to substantiate the normativity thesis ultimately appeals to our being guided by – and so following – internalised rules, in order to fix the meaning of our utterances and the content of our thoughts. And so, it naturally raises the question: what is involved when we do so; when we follow such rules?

As we have previously noted, it seems uncontentious that following a rule cannot consist merely in acting in accordance with a rule.[7] Indeed, we can easily imagine cases in which an agent conforms to a rule by accident, without having the rule in mind – or even, without being familiar with the rule. In jotting numbers down on a piece of paper – or in ordering them in his mind –, an individual, otherwise uncultivated in mathematics and thus ignorant, for example, of the rule of formation of the Fibonacci sequence, could nonetheless happen to produce the beginning of such a sequence.[8] Likewise, while playing on an empty board, a child might move a rook according to the rules of chess even though they have never learned how to play chess and how to move the different pieces. We can also imagine cases in which an agent conforms to a rule, not because they are committed to conform to the rule, but because they have been submitted to some kind of brute

[7] I am using 'acting' in the broadest possible sense to encompass both the production of overt behaviour (including linguistic behaviour) and the entertainment of thoughts (including the transition from one thought to another).
[8] The example is from Bridges 2014.

conditioning. In all these scenarios, the agent does act in accordance with the rule, but they do so without understanding, and *a fortirori* without any consideration for the requirements of the rule (i.e. for the normative constraints the rule imposes). They are ultimately not responsive (or so it seems) to a norm, even if they happen to satisfy its demands. Of course, in such scenarios, the behaviour of the agent could display certain regularities, but even these regularities are not sufficient to constitute their following a rule. The point is that even if the agent were to conform again and again to a rule – even if they came to exhibit a disposition to do so –, they might still not be acting *in light* of the rule. After all, in the scenarios in which the agent is submitted to brute conditioning, it is not clear how they would be different from an automaton, or from a parrot mindlessly responding to features of its environment.

What, then, is needed for an agent to genuinely be following a rule? A common answer is that the agent must take the rule as a *reason* for acting as they do.[9] The rule must not only explain their behaviour, it must also be part of their reason for acting. That is to say, following a rule *R* consists in invoking, if only implicitly, *R* in contexts of practical reasoning. A rule *R* would not be followed by an agent if it were not part of their resolution to act in one way rather than another; if it did not bring about their action. It implies, ultimately, that following a rule *R* is not possible without some kind of commitment to conform one's behaviour to *R* – which is to say that following a rule requires an *intention* to do so. Thus, according to this line of thought, there is, we could say, an intentional requirement on following a rule that distinguishes it from merely conforming to a rule.

[9] See for example Boghossian 2005b, 2008 and 2012, Glüer and Pagin 1999, Glüer and Wikforss 2009, Pettit 2002 and 2016, Wright 2007 and 2012.

Such a requirement is appealing, I think, for two reasons. First, it clearly explains why the scenarios previously imagined are not cases of rule-following. We were inclined to say that in fortuitously writing down the first few numbers of the Fibonacci sequence, an individual was not following a rule. The intentional requirement substantiates that insight: if it is not a case of rule-following, it is because it is done without any intention to conform to the rule of formation of the Fibonacci sequence. The same is true of the example of the child who moves a rook on a chessboard: if, in doing so, they are not following the rules of chess, it is because they do not, at any point, *intend* to abide by these rules, being, by stipulation, ignorant of them. As for an agent who would exhibit a disposition to abide by some rule: if the disposition is not brought about by the agent's internalised commitment to the rule, but solely by the brute conditioning to which they have been subjected, then it is not a case of following a rule. It is nothing more than a regularity in behaviour.

Second, and more importantly, the intentional requirement makes room for the idea that our patterns of thought and behaviour can be justified, that they are not mere 'leaps in the dark'. And even more so, that they can be justified in light of our own commitments. Nothing forces me to abide by the rules of chess. After all, I might decide to do any number of things with chess pieces and a board (e.g. play a new game I have invented, depict the dispositions of troops on a battlefield, etc.). However, if I *intend* to abide by the rules of chess, then the rules give me reason to move my rook in straight lines or, in certain circumstances, to castle with the king, etc.

Yet, despite these merits, many authors are of the opinion that such a line of thought faces important difficulties when we assume that meaning and content are themselves determined by rules. Indeed, for them, it is difficult to see how the intentional requirement

could then be satisfied (see most notably Boghossian 2005b, 2008 and 2012 and Glüer and Wikforss 2009). The general idea is that in trying to satisfy the requirement, we would inevitably embark on an endless regress. For Boghossian, it is rather straightforward: intending to conform to a rule already requires if not the ability to speak meaningfully, at least the ability to entertain contentful thoughts. Yet, under the assumption, thinking is a matter of following a rule, which requires, in turn, an intention to conform to a rule, etc. Glüer and Wikforss call it the 'regress of contents'; the moral being that no meaningful utterances or contentful thoughts – and as such no intention – can be invoked as a condition on following a rule, if the very possibility of enjoying an item as meaningful or contentful depends upon our ability to follow a rule (see Glüer and Wikforss, 2009, 57).

In light of our discussion of the regress of interpretations (see *supra* 1.2), this is hardly surprising. To assume that a sign's meaning is determined by a rule implies that the sign, in itself, just 'stands there', that something else must give it life. And this is precisely the picture of sign – as essentially 'dead' or 'normatively inert' – that gets the regress of interpretations going. (Or so we have argued.) To be sure, the intentional requirement forces upon us the same picture: it assumes that a sign – or even an object, like a chess piece – just stands there, silent, until our commitment to a rule brings into existence various normative constraints. The regress' endlessness is made salient by the underlying thought that the rule to which one commits and that determines the sign's meaning is, itself, a further sign That is, by the underlying thought that the rule is always an *expression* of the rule. If this is right, then, the 'regress of content' is really just another name for the philosophical perplexity that is voiced by the regress of interpretations.

What about the two other regresses discussed in relation to the intentional requirement: the "minor premise regress" and the "inference regress" (see Boghossian 2008 and 2012 and Wright 2012)? As we have seen, the intentional requirement suggests that following a rule *R* depends on one's internalisation of the rule *R*. But for many it seems clear that the rule, in itself, will not guide one's performances. As Wright puts it: "one's knowledge of the rule is but one ingredient in one's movement to a correct application of the rule" (Wright, 2007, 491). I can intend to conform to a rule of the form "if conditions C obtain, then do A!". However, the rule will only guide my performance when in conjunction with the belief that conditions C obtain. Hence the suggestion that following a rule *R* is ultimately a case of practical reasoning involving *R*. Schematically, it gives us something like:

(*R*): If conditions C obtain, then do A!

(*P*): Conditions C obtain.

(*C*) Do A!

The problem is that, if rules determine meaning and content, then the content of the minor premise "conditions C obtain" would depend upon further rules, whose potency would, in turn, depend upon the recognition that some other conditions C obtain, etc. Wright calls it the "minor premise regress". It is also discussed, at length, by Glüer and Wikforss. They invoke it against the normativity of belief, but it is clear that, for them – as for Wright –, it equally applies to meaning and content generally (see Glüer and Wikforss, 2009, 57, note 58). Accepting both the intentional requirement on following a rule and the assumption that rules determine meaning and content appears, once again, untenable.

In light of these considerations, Glüer and Wikforss conclude that the only way to defend the claim that meaning and content are determined by rules would be to abandon the intentional requirement – and so to recognise that rules could be followed without any commitment to them. However, they find this hard to accept. As they put it:

> Given how overwhelmingly intuitive the idea of an intentional [requirement] on rule-following is the choice might well seem to be between Scylla and Charybdis. For how could there even be any distinction between following a rule *R* and merely regularly according with *R*, if the intentional [requirement] goes by the board? (Glüer and Wikforss, 2009, 57)

For them, making sense of the normativity thesis in terms of one's commitment to act in accordance with a rule – i.e. in terms of one following a rule – is thus either fraught with intractable difficulties or nothing more than an "idle label" (Glüer and Wikforss, 2009, 59). In either case, it appears to be condemned to failure.

Interestingly, while they also recognise the intuitive appeal of the intentional requirement, Boghossian and Wright are a bit more cautious. The issue, for them, is that the requirement leads to another regress (the 'inference regress'), even if we abandon the normativity thesis and concede that mental contents – or linguistic meanings – are not, in general, determined by the rules one is following. Indeed, given the intentional requirement, following a rule *R* involves some kind of inference from the rule *R* and a further premise. The problem, for Boghossian and Wright, is that inference is clearly a case of rule-following. To follow any rule *R*, one would thus need to follow some rule of inference (in this case, potentially a rule like "From 'If C, then A!' and 'C', conclude 'do A!'"). And, to follow this rule, one would require yet another inference, etc. For them, the regress is not only vicious, but unavoidable so long as we maintain an intentional requirement on rule-following. Boghossian and Wright's conclusion is that the intentional

view on rule-following must be amended. Appealing as it is, the view, in its actual formulation, reveals itself to be self-defeating under closer scrutiny.

The 'minor premise regress' and the 'inference regress' emerge as soon as we take cases of rule-following to be dependent upon an act of practical reasoning. As I see it, the apparent inescapability of this thought ultimately finds its root in the picture of meaningfulness – or contentfulness – which underlies both the intentional view on rule-following and the insight that meanings – and contents – are determined by rules. This is what gets the regresses going: the thought that a sign could not *tell* an agent what to do, were it not for the agent's prior commitment to – or enjoyment of – the contentful expression of a rule.

Our earlier discussion of McDowell's view on following a rule is, on this point, informative (see *supra* 1.2): if we acknowledge, like McDowell, that meaningful signs, in themselves, can guide us, then there is never any need to interpose a rule between the sign and the normative constraints it is submitted to. Our following a signpost, for example, does not need to depend on our commitment to – or awareness of – any rule. But then, if no rule is involved in our following the signpost, we lose the major premise and with it, the need for a minor premise. And if there is no major premise nor minor premise, then no deductive inference is needed anymore.

If this is right, then both regresses are really just further variations on the regress of interpretations. Boghossian and Wright would thus be erring in suggesting that the 'inference regress' is a further challenge *over and above* the minor premise regress or the regress of content. [10] As would Glüer and Wikforss in their diagnosis of these regresses as

[10] I take this to be very much in line with Miller's view on both regresses (see Miller 2015).

emerging from the *interaction* of the intentional view and the insight that meanings – and contents – are determined by rules.

Why does this matter? Well, in our view, and in contrast with Glüer and Wikforss's suggestion, there would be little hope to substantiate the normativity thesis by appealing to meaning-determining – or content-determining – norms *even if* we were to abandon the intentional requirement. The issue with the second strategy to substantiate the normativity thesis is more fundamental.

5.4 Ginsborg's Primitive Normativity

It is in light of the apparent failure of the second strategy that Ginsborg's defence of the normativity thesis – and account of meaning and rule-following – emerges as a promising contender. After all, it precisely rests on the insight that there is a way to follow a rule prior to and independent of the grasp of any rule.

Ginsborg starts by suggesting that what crucially distinguishes following a rule from merely acting in accordance with a rule is that the former presupposes understanding of what is appropriate, at each step, while the latter does not (see Ginsborg, 2012, 134). The idea is that an agent can carry out certain tasks – in response to an order for example – without really understanding what is asked of him. In such a case, the agent conceivably acts in such a way as to conform to a rule, but, for Ginsborg, we would be reluctant – and rightly so – to say that he is following a rule in any way. He would ultimately be lacking the consciousness that what he is doing is appropriate, wandering around without a genuine appreciation that his actions accord with the order that he was given. Ginsborg uses the language-game of the builder and the assistant to illustrate her insight. The assistant might

well be disposed to bring a slab to the builder whenever the builder utters "Slab!" – and so be disposed to act in accordance with a certain rule determining how it would be correct to respond to the order "Slab!". Yet – and as we have repeatedly argued – the assistant understanding what is being asked of him cannot be reduced to the possession of such a disposition. After all, he could simply be reacting "blindly" or in a "jack-in-the-box" manner to the orders of the builder (see Ginsborg, 2012, 134). The point is now familiar: a regularity in behaviour (e.g. bringing a slab to the builder whenever one hears the order "Slab!") is not sufficient to constitute an agent's following a rule since the behaviour of an automaton – clearly lacking understanding, and *a fortiori* understanding of any normative constraint – can still exhibit this kind of regularity. But, for Ginsborg, it would be a different story if, in addition to predictably bringing a slab to the builder when receiving the order "Slab!", the assistant was also disposed to see his responses to that order as *appropriate*. Bringing a slab to the builder can be done with or without understanding, but it is only in the former case – that is, when a consciousness of appropriateness accompanies the assistant's actions – that we can legitimately talk of rule-following.

Now Ginsborg insists that this consciousness of appropriateness is independent of the agent's "taking what [he] is doing to accord with a rule" (see Ginsborg, 2011, 235). That is to say, for her, the assistant can ultimately take his response to the order "Slab!" as appropriate – or correct –, even though he does not recognise, at any point, that what he is doing satisfies the requirements of a rule. Of course, it is conceivable that the assistant sometimes recalls a rule before complying with the builder's order. For example, at some point, he might have memorised a chart – or written directives – and might subsequently refer to it when he has doubts as to which building block to bring back to the builder. I do

not think Ginsborg wants to deny such possibilities. As I take it, her point is rather that such a way of 'following a rule' is secondary. Some of our everyday activities may indeed be governed by rules in that way, but, for Ginsborg, this observation should not obscure the fact that it is possible to take one's own action as appropriate *simpliciter* – and so, to follow a rule without any reference to an "antecedently recognized rule" (Ginsborg, 2011, 233). Ultimately, it is *because* we have a primitive consciousness of appropriateness that explicit rules can come to play the role they do.

For Ginsborg, these considerations on rule-following also apply to meaning. The idea is that the distinction between the meaningful employment of a linguistic expression and the mere production of noise can be drawn along the same lines as the distinction between following a rule and merely conforming to a rule. Clearly most competent speakers never learn any explicit rules as to how to use the colour word "green". There is thus no sense in saying that such rules guide their uses of "green". But for Ginsborg, they are still employing the word "green" meaningfully if, whenever they do so, they have the consciousness of employing it appropriately. Nothing more is needed to distinguish competent speakers from the parrot trained to utter the sound "green" whenever it sees something green. If it has been properly trained, the parrot will, time and again, conform to certain rules about how to use "green", but it will ultimately lack the consciousness of going on appropriately. And for Ginsborg, this is sufficient to deny the parrot the status of speaker – and so of rule-follower (as far as the use of "green" is concerned, of course). The same applies in the case of the enjoyment of contentful mental items.

Hence, for Ginsborg, rule-following – including meaning and content – does not ultimately depend upon the commitment to – nor even upon the awareness of – a rule. It

only calls for a "primitive normativity", for a normativity that can manifest itself prior to the establishment of any rule (see Ginsborg, 2011, 233).[11] That is what she is suggesting, in a nutshell, when she talks of a sense of appropriateness *simpliciter*. Ginsborg thus clearly disagrees with the proponents of the intentional view that what is present when an agent is following a rule, but absent when he merely acts in accordance with a rule, is to be explained in terms of commitment to conform to a rule. As a matter of fact, she even explicitly criticises recent attempts to do so (see Ginsborg, 2011b, 233 and Ginsborg, 2012, 137, respectively).

In insisting that rule-following is a matter of understanding what is appropriate – and so that it involves a primitive normativity –, she is advocating for an alternative to the intentional view and its underlying picture. In parallel, she is proposing a new way to substantiate the normativity thesis, one that does not appeal to meaning-determining – or content-determining – rules. Of course, we have not yet settled the question as to whether or not her account really amounts to a viable alternative to the intentional view and if it can really substantiate the normativity thesis – it will be the focus of the next section. But a case can already be made that, at the very least, it avoids the regress of interpretations – and its three variations.

Indeed, for Ginsborg, one's responsiveness to a sign, in cases of rule-following, is not dependent upon one's enjoyment of a further sign (i.e. a rule). And so, the question of how that further sign could ever be contentful never arises. [12] The regress of content never

[11] The sense of appropriateness proposed by Ginsborg is also 'primitive' in the sense that it is constitutive of meaning without being essentially semantic – nor even intentional (see for example Ginsborg 2012 and 2020).

[12] It could be argued that Ginsborg merely pushes back the difficulty and remains, as such, prisoner of the regress of contents – and so of the regress of interpretations. (This has been suggested to me by Anandi Hattiangadi in a private conversation.) After all, in her account, rule-following still depends upon the enjoyment of a particular kind of mental item (i.e. an awareness of primitive appropriateness). And to some

gets going. Relatedly, for Ginsborg, cases of rule-following do not depend upon the drawing of inference – involving further signs –, thus avoiding both the minor premise regress and the inference regress.

5.5 Guidance, Objectivity and Epiphenomenal Understanding

However, Ginsborg's account of meaning and rule-following also raises certain worries (for recent discussions of some of these worries see Haddock 2012, Miller 2017 and Verheggen 2015). In what follows, I will focus on three of them, which I take to be the most serious.

First, there is an issue with her insistence that the understanding that is present in cases of rule-following but absent in the cases of mere conformity to a rule ultimately amounts to a consciousness of appropriateness that *precedes* not only one's grasp of a rule, but also one's ability to discriminate between responses that are appropriate and responses that are inappropriate. As she puts it, discussing the language-game of the builder and the assistant:

> When the assistant takes his bringing of a slab to be appropriate to the builder's shout of 'Slab!', his attitude is not that other responses would or might be inappropriate or wrong, but simply that they would lack the quality of appropriateness which belongs, under those circumstances, to his bringing a slab. (Ginsborg, 2012, 139)

Thus, the assistant could respond with understanding to the order "Slab!" and yet be unable to see his bringing a beam or a pillar as inappropriate. Ginsborg sees this concession as crucial to her account. Short of it, the assistant's consciousness of responding appropriately

---

it might seem enough to get the regress going. However, this is only the case if the posited primitive awareness has *intentional* content. I take that her recent talk of a primitive "going on" rather than a primitive "awareness of appropriateness" is an attempt to diffuse such concerns. (see for example Ginsborg 2020 and 2021).

to the order "Slab!" by bringing a slab to the builder could not be constitutive of his understanding "Slab!".

> For in order for [the assistant] to suppose that a response could be inappropriate, as opposed to merely lacking appropriateness, he must already regard 'slab' as a meaningful term and not just […] 'mere noise'. And that is something he can do, on the proposal […], only if he can invoke the sense of appropriateness under discussion here. (Ginsborg, 2012, 139)

The problem with such a concession is that it is not at all clear how one then gets from the distinction between appropriateness and lacking appropriateness to the distinction between appropriateness, inappropriateness, and lacking appropriateness. How can the class of actions initially seen as lacking appropriateness be later divided into those that are, as a matter of fact, inappropriate (e.g. bringing a beam or a pillar instead of a slab) and those that really merely lack appropriateness (e.g. adjusting one's hat while carrying a slab)?[13] If there is nothing more to the order "Slab!" being endowed with meaning than this primitive consciousness of appropriateness, what grounds the fact that bringing a building block other than a slab is inappropriate? Ginsborg remains regrettably silent on this point.

Ginsborg is wary of talks of guidance in discussing rule-following activities. But if to follow a rule does not require to be guided by the rule (i.e. by a propositionally articulated expression of the rule), it requires at the very least a kind of sensitivity to the requirements of the rule. The assistant might not have to grasp any rule in order to follow the orders of the builder. He might not have to think to himself, at any point "if the builder shouts "Slab!", bring a slab!", nor anything of the sort. Still, it seems the assistant must be in a position to recognise not only that certain actions satisfy the builder's order, but also that others contravene to it. Bringing a beam or a pillar instead of a slab is not irrelevant to

[13] Miller raises essentially the same question (see Miller, 2017).

the builder's order; it goes against it. And seeing this difference – like seeing what the order calls for – is, I would contend, constitutive of the assistant's understanding of the order. Short of this ability to discriminate between those actions that are appropriate, those that are inappropriate and those that simply lack appropriateness, it is difficult to see how the assistant would be different from his counterpart who, after having been submitted to brute conditioning, simply acts in a jack-in-the-box manner. That is to say: it is difficult to see how the assistant's consciousness of appropriateness would amount to a genuinely normative attitude – and *a fortiori* to a genuine form of understanding – rather than to a mere feeling of satisfaction emerging from conditioning.

To be responsive to the requirements of a rule is to be in a position to appreciate different actions as having different normative values. It is to be able to see that certain actions satisfy the requirements of the rule, that others violate them, while others do neither. Ginsborg introduces the notion of primitive normativity, with the promise that it can account for that responsiveness. However, it does not seem as if the consciousness of appropriateness she posits can really do so.

In contrast, our own proposal has no particular difficulty in accommodating such distinctions. Indeed, under our proposal, whether an action is *appropriate* in response to the order "Slab!" – or as we have put it whether an action fits the order "Slab!" – is grounded on the very meaning of the order. And, contrary to Ginsborg's consciousness of appropriateness, the meaning of the order does *polarise* the world, does discriminate between responses that are fitting and responses that are *un*fitting. That is what meanings do! If the assistant understands – or grasps – the meaning of the order "Slab!", then he

understands that his bringing a slab fits the order (see *supra* 4.3). But he also understands that his bringing a pillar or a beam is *un*fitting, that it goes *against* the order.

Here is another way to look at it: in suggesting that certain responses might be more *appropriate* or *fitting* than others (see *supra* 4.2.5), we implicitly acknowledged that there was a 'hierarchy of fittingness' among the possible responses of the assistant. But clearly, this hierarchy spans from responses that are altogether *un*fitting to those that *best* fit the meaning – and more precisely, in our view, the purpose – of the builder's order.

This 'hierarchy of fittingness' also highlights how, under our proposal, certain actions might be neither fitting nor unfitting. Think again of the imagined case in which the assistant turns on himself three times before bringing a slab to the builder. We have suggested that such a response would be, in most circumstances, *less* fitting than directly bringing a slab to the builder. That is, in most circumstances, turning on oneself three times would not be *irrelevant* to the overall fittingness of the assistant's response. But adjusting one's hat, while bringing a slab to the builder might well be *irrelevant*, that is, without effect, on the overall fittingness of the assistant's response.

Ginsborg's account faces a further potential difficulty when it comes to the objectivity condition on rule-following. As we have repeatedly argued, following a rule might depend on one's responsiveness to the requirements of a rule, but these requirements cannot then collapse with what one takes them to be (see for example *supra* 1.4 and 4.2.1-4.2.2). What is, as a matter of fact, appropriate in light of the requirements of a rule cannot simply be what one thinks is appropriate. Otherwise, we could not talk of anything being appropriate. Clearly, we can be mistaken in following a rule without realising it. And to make sense of such mistakes there must be something that *grounds* the requirements of a

rule and which is independent of – or at least not constitutively dependent on – what we take these requirements to be. That bringing a slab in response to the order "Slab!" is appropriate cannot merely consist in the assistant taking it to be appropriate. Whether or not his response is appropriate must depend on something over and above that very response. The point evidently applies to language use and to thought more generally.

As we have argued at length, our proposal meets the objectivity condition on rule-following: in our view, the normative constraints to which a sign is submitted – what we call here "the requirements of a rule" – are grounded on the meaning of the sign. And that meaning, even if it is a creature of ours, remains part of the fabric of the world (see *supra* 4.2.2).

The question is whether or not Ginsborg's account can also meet the objectivity condition. At first sight, one might be doubtful. After all, in explaining the assistant's understanding of the order "Slab!" solely in terms of his disposition to respond in a certain way and of his taking his response to be appropriate, Ginsborg gives herself little room to manoeuvre. The assistant's disposition to bring a slab whenever he hears the order "Slab!" cannot ground what is, as a matter of fact, an appropriate response to that order. As Kripke famously remarks, the assistant's disposition has implications for how he will respond to the builder's order, but not for how he ought to. After all, the assistant could also have a disposition to make mistakes when certain conditions obtain (e.g. when there is too much noise around, when the builder gives the order without looking in his direction, etc.).

Not any disposition would do, but what about the disposition that encompasses the assistant's responses *under ideal conditions*? The conditions to which the assistant is actually submitted are not always ideal and so it is conceivable that his actual responses

would sometimes deviate from the responses he would be ideally disposed to produce. As such, what is the appropriate response to the builder's order would not collapse with what the assistant takes it to be in any given case. The proposal might be appealing at first sight, but it does not sustain closer scrutiny. The main problem is that we cannot specify these ideal conditions without already presupposing what they are meant to establish, namely the set of appropriate responses to the builder's order. Imagine that the assistant has a disposition to bring a pillar when he cannot properly hear the orders addressed to him by the builder (he does not necessarily believe that the builder shouted "Pillar!" in each case; often, he only guesses). As such, there are occasions in which the builder shouts "Slab!" and the assistant brings him a pillar instead of a slab, because the surrounding noise prevents the assistant from hearing the order properly. How are we to decide that these conditions are interfering with the assistant's disposition to conform to the builder's order? That is to say, how are we to decide that these conditions are not part of the ideal conditions? The only way, it seems, would be to say that in bringing a pillar, he is not conforming to the order of the builder. He is not responding appropriately. But that is precisely what needs to be established…

We are thus left with the assistant consciousness of appropriateness, which does seem insufficient to make sense of the distinction between a response being appropriate and a response seeming appropriate, without being so. Ginsborg tries to address this problem, but her response appears unconvincing.

She brings the problem into focus by discussing the example of an 'abnormal' child who does not continue the series '+2' like we do and claims that "43" – rather than "42" – comes after "40". Imagine furthermore that the abnormal child takes "43" to be the

appropriate continuation of "40", in the same way that a normal child would take "42" to be the appropriate continuation. Both could be equally convinced that their response is the proper way to continue the series. So what makes it the case that only the latter's response is appropriate – or as we would put it: fitting? Ginsborg does want to preserve the insight that the abnormal child's response is a mistake, not only for us, from our point of view, but as such, given the requirements of the rule. However, she offers little to substantiate the compatibility of her account with that insight. As far as I know, her only defence pivots around the suggestion that: "while we might have the theoretically motivated thought that the two children's claims are equally legitimate, that thought is trumped by the pretheoretical intuition that "42" is appropriate and that "43" is not" (Ginsborg, 2011b, 250). But this appeal to our "pretheoretical intuition" must appear unsatisfying. It ultimately remains silent as to what in her account of rule-following satisfies the objectivity condition.

In contrast, our proposal has no difficulty in accounting for the insight that the abnormal child's response is a mistake. It even offers us the conceptual apparatus necessary to formulate a hypothesis as to why the abnormal child makes such a mistake. Indeed, under our proposal, that "42" is the right continuation is grounded on the basic pattern instantiated by the beginning of the series "+2". To the question: "why is "42" the correct continuation?", we can thus answer: *because* it is the answer that fits the beginning of the series. And it fits the beginning of the series, not because of any intuition we have on the matter, but because of how things are. That is, because of the basic pattern that is, as a matter of fact, instantiated by the beginning of the series. And so, if the abnormal child repeatedly fails to continue the series with "42", it might be, for example, because he is

blind to that basic pattern. His constitution might be such that is simply not *sensitive* to that basic pattern – and so, to that meaning.

There is one last worry that needs to be addressed. Ginsborg is keen to describe the assistant bringing a slab in response to the order "Slab!" and seeing his response as appropriate as acting "with understanding, or intelligently" (Ginsborg, 2012, 128). This is of course unsurprising: as we have seen, this is how she substantiates the insight that there is something fundamentally different between what the assistant is doing and what an automaton would be doing, if it were devised to bring a slab whenever it hears the sound "Slab!". However, she also takes seriously the claim that the pattern of responses of the assistant is dispositional and as such, potentially indistinguishable from that of such an automaton. It is, as a matter of fact, her starting point in laying down her account of rule-following (see Ginsborg, 2012, 133-134).

The worry goes like this: in her view, the assistant's understanding amounts to a consciousness of appropriateness. But such a "consciousness" – she also sometimes talks of a "sense", "attitude" or "awareness" – does not seem to play any role in *bringing about* the responses of the assistant. That is, in her view, the assistant's understanding of the requirements of the rule appears to parallel what he does, without, however having any repercussions on it (see Ginsborg, 2011b, 238 and Ginsborg, 2012, 137-138). She seems to be suggesting essentially the same thing when she discusses why a child learning a language is unlike a parrot being trained to utter mere sounds. She concedes that the child and the parrot could be disposed to utter the same sound ("green" for example) under the same conditions (when a green object comes into view). But then adds: "the situation of the child differs from that of the parrot in that the former takes herself […] to be responding

appropriately" (Ginsborg, 2011, 237). But then, if this is indeed her view and it is the *only* difference, the child's understanding – just like the assistant's – would be an epiphenomenon.

I would contend that this is incompatible with the guidance condition. For the assistant to be guided by the order "Slab!", it does not suffice for him to act *with* understanding. He must act *on* that understanding. And this implies precisely that his understanding is involved – and even causally involved – in bringing about his response to the order "Slab!". The point, I take it, is a general one: it holds however we flesh out the assistant's understanding.

The first issue we raised against Ginsborg's account, was that if the understanding that is constitutively present in cases of rule-following is nothing more than the agent awareness of the (primitive) appropriateness of their action, then "understanding" might be a misnomer. The posited awareness simply does not seem to amount to a genuine understanding. The worry now is that, even if it were to amount to a genuine understanding, it would still not be 'connected' in the right way to the agent's action. (As a matter of fact, it would not be connected at all.)

To be sure, our proposal does not raise any such worry, as it explicitly stipulates that an action performed by an agent does not count as a case of rule-following if the agent's understanding (of the relevant relations of fit) does not play any role in bringing about their action. For us, the assistant's response to the builder's order is only a case of rule-following if the order is causally involved in bringing about the assistant's response, *and* if that causal involvement partly rests on the assistant's understanding (of the relation of fit that holds

between the order and his action).[14] Under our proposal, there is thus a fundamental difference between *how* the builder's order – causally – brings about the assistant's response and *how* it – causally – brings about the response of the automaton.[15]

5.6 Intelligent Action

There is a further point worth discussing: as we have mentioned, Ginsborg says of the assistant that, contrary to the automaton, he acts "with understanding, or intelligently" when he responds to the builder's order. But, should we really consider "acting with understanding" and "acting intelligently" to be synonymous? I think there might indeed be an important connection between understanding and intelligent action, but it is not as straightforward.

For Ginsborg, the assistant's response to the builder's order might be – behaviourally – indistinguishable from that of the automaton. Just like the child's use of "green" might be indistinguishable from the pattern of use of the parrot. Now, I think both claims are perfectly reasonable, as illustrations of the idea that one's conformity to normative constraints cannot guarantee that one is indeed acting *in light of* these constraints. That is, they are perfectly reasonable as expressions of a *conceptual possibility*: there is nothing contradictory in the thought that a being acting in a jack-in-the-box manner and a being acting in light of normative constraints both conform to said normative constraints, and so act in the same way. But, I would contend that the claims about the builder and the child are rather implausible as hypotheses about the kind of pattern of behaviour the assistant

[14] See *supra* 4.3.1
[15] That is, of course, assuming that, unlike the assistant, the automaton is not sensitive to the meaning of the order, and so does not understand the relation of fit that holds between the order and its response.

and the child would likely engage in *if* they were indeed sensitive – and ultimately responsive – to the relevant normative constraints.

Our proposal enables us to make that point quite clearly. Let us assume that the assistant is indeed sensitive to the meaning – and more precisely, to the purpose – of the builder's order, and, as such, understands the various relations of fit that hold between the order and possible responses to it. Let's further assume that the assistant's understanding of these relations of fit – and so, of these normative constraints – is causally involved in bringing about his responses to the builder's order. Now, if that is the case, why suppose that the pattern of his responses would be indistinguishable from that of a being – like the automaton – which, by stipulation, lacks such understanding?

In our view, both the assistant's and the automaton's responses are, in some way, causally connected to the builder's order. But the assistant and the automaton are not sensitive and ultimately responsive to the same features of the builder's order. So it seems rather plausible that the order would not always bring about identical responses from both of them.

Why does this matter? Well, if understanding does not run parallel to action, as an epiphenomenon, if it makes a difference in how one acts, then there seems to be room for the thought that understanding could *manifest itself* in action. The thought, I think, is quite appealing – and finds expressions in the writings of many authors (see for example Baker and Hacker 1984, Brandom 1998 and more recently Connelly 2012).[16] The difficulty is to

[16]. Interestingly enough, such a thought may also find expression in Descartes's writings. In a passage of the *Discourse on Method* surprising germane to our present concerns, he asks how we can distinguish a human being from an automaton or a parrot. His answer goes as follows: he concedes that we could create machines that "do many things as well as, or perhaps even better than, men". However, he insists that such machines "would infallibly fail in certain others, by which we would discover that they did not act by understanding, but only by the disposition of their organs" (Descartes, 1960, 42). He argues that the same is true in the case

reconcile it with the *conceptual possibility* just highlighted. How can we make sense of the agent's understanding *showing itself* in how he acts, if we recognise that there is no contradiction in thinking of the agent as behaviourally indistinguishable from an automaton?

In my opinion, this is only possible if we conceive of the agent as acting not merely *with* understanding, like in Ginsborg's account, but *on* that understanding, like in our proposal. The actions of the agent cannot be a *criterion* for understanding. If we are not to lose the conceptual possibility of the behaviourally indistinguishable automaton, understanding the order of the builder cannot *consist in* responding to that order in a particular manner. So then, the only way in which the agent's understanding could *manifest itself* in the agent's actions is if it is – causally – involved in bringing about these actions.

As I see it, the kind of actions that are brought about on the ground of the agent's understanding could aptly be called "intelligent actions". And, in many cases, such actions might well be – behaviourally – different from "mechanical actions" that are performed in a jack-in-a-box manner (i.e. in the absence of understanding).[17] It might just be a matter of fact that certain actions are only performed in response to particular signs, objects or situations by beings who are capable of understanding – and so, in our view, by beings who are appropriately sensitive to the meaning of these signs, objects or situations. Ginsborg rightly points at a connection between understanding and intelligent action; she just does not, in my opinion, flesh it out in the right way.

---

of the parrot. Thus, for Descartes, that the automaton and the parrot do not possess reason reveals itself in how they act.

[17] I cannot elaborate on the matter here, but I take Wittgenstein to be suggesting something along those lines in a few remarks of the *Zettel* in which he comes back to the language-game of the builder and the assistant (see *Z* §§ 98-108).

Conclusion

As our work comes to a close – though certainly not our thinking on the topic –, one may wonder why we proceeded as we did. If our work was indeed an attempt at constructive philosophy – even of the unproblematic kind we have tried to circumscribe – , it might seem perplexing that we have spent so much time discussing the views of McDowell, Merleau-Ponty and Wittgenstein. Indeed, if our aim was ultimately to lay the groundwork for a philosophically substantive account of rule-following, our discussion of their views might look like an unnecessary exegetical detour.

Now, I certainly do not want to deny that there was an exegetical dimension to our work. But I would like to suggest that it was, as it were, secondary. Here is one way to look at it: these authors and their views were not for us an object of exegetical curiosity. McDowell, Merleau-Ponty and Wittgenstein were rather like *interlocutors* and our critical engagement with their views, akin to a *dialogue*. So, if our work reveals a certain exegetical rigour, it is not just for the sake of it. It manifests the kind of respect we think is owed to one's interlocutors – even interlocutors from the past.

Our focus on the primary texts of McDowell, Merleau-Ponty and Wittgenstein – rather than on the secondary literature they have given rise to – should be seen in that light. As should our attention to the details, not only of *what* they say, but also of *how* they say it. To treat them as interlocutors is to take the specificities of their philosophical voice seriously.

So then, through our dialogue with McDowell, Merleau-Ponty and Wittgenstein, we have tried to sketch an answer to two questions: (i) which feature of signs are we *responsive to* in cases of rule-following? and (ii) what is the relation between these signs and our responses to them in such cases? Evidently, the proposed answer incorporates many insights we have attributed to either McDowell, Merleau-Ponty or Wittgenstein, but I ultimately take it to be our own.

After all, in fleshing it out, we have not hesitated to distance ourselves from some of their respective views. For example, we have explicitly disagreed with McDowell on where to locate the intentional bedrock that grounds meaning, understanding and rule-following. We have also distanced ourselves from his – and Wittgenstein's – quietist approach to the philosophical perplexities surrounding rule-following. We have also been at times critical of Merleau-Ponty, not shying away from highlighting some persisting imprecisions and internal tensions in his conception of motivation.

Still, despite these various points of departure, I would contend that we would not – and could not – have fleshed out our proposal independently of our critical engagement with the views of McDowell, Merleau-Ponty and Wittgenstein. That is, our proposal would not – and could not – have emerged outside the distinctive conceptual space that arises from the meeting of their respective views. If this is right, then our apparent exegetical detour was not really a detour at all; there was simply no road to our proposal but a meandrous one.

We have been quite transparent, throughout the present work, about its preliminary nature. By our own admission, it is at best the groundwork for a substantive account of rule-following. Consequently, we have been keen to highlight, along the way, many lines

of inquiry that remained to be pursued. Instead of coming back to these, I would like to conclude by briefly sketching a further possible line of inquiry – one we have not yet discussed, and which emerges from our discussion of the relation between understanding and intelligent action.

At the end of the last chapter, we suggested the following: to the extent that our responses to signs, objects and situations are grounded on our understanding of the relations of fit these signs, objects and situations are involved in, they might be – behaviourally distinct – from the responses of beings, like the automaton or the parrot, who do not possess such an understanding. Some patterns of behaviour we engage in might simply be, as a matter of fact, out of reach of these beings. We then suggested that there might be a helpful – and substantive – distinction to be drawn between intelligent and – merely – mechanical actions. Whereas we would, by large, engage in the former, beings like the automaton or the parrot would be confined to the latter.

But then, we have also suggested the following: one's understanding of the relations of fit signs, objects or situations are involved in is grounded on one being appropriately sensitive to their meanings – which, in turn, is dependent not only upon one being initiated into the relevant practices, but on one actively keeping track of these meanings. This leads us to the following thought: our engaging in intelligent actions is dependent upon our having meanings in view. And having meanings in view is not given to us, once and for all. As individuals, we may lose track of meanings – fail to keep up with our creatures, as we once framed it. At times, we may also satisfy ourselves with a confused view on meanings, be contented with our poor acuity – to the point of blindness. As a culture, we

may fail to initiate future generations into the relevant practices, thus progressively loosing track of meanings that were once collectively in clear view.

In its most tragic formulation, the thought is this: our having meanings in view is a fragile achievement. It would not take much for us – either individually or collectively – to lose it and to end up confined, like the automaton or the parrot, to mechanical gestures and lifeless utterances. We are, we could say, always in danger of becoming a parody of ourselves. In light of this, we may thus wonder: do we have a responsibility to keep meanings in view?

Bibliography

Ayer, A. J. (1945). The Terminology of Sense-Data. *Mind*, *54* (216), 289-312.

Baker, G. (2004). *Wittgenstein's Method*. Oxford: Blackwell Publishing.

Baker, G. & Hacker, P.M.S. (1984). On Misunderstanding Wittgenstein: Kripke's Private Language Argument. *Synthese*, *58*, 407-450.

Baker, G. & Hacker, P.M.S. (1985). *Wittgenstein: Rules, Grammar and Necessity*. Oxford: Basil Blackwell.

Baker, G. & Hacker, P. M. S. (2005) *Wittgenstein: Understanding and Meaning*, 2nd revised edition, Oxford: Blackwell.

Beaney, M. (2006). Wittgenstein on Language, in E. Lepore and Barry C. Smith (eds.) *Oxford Handbook to the Philosophy of Language*. Oxford: Oxford University Press.

Brandom, R. (1998). *Making it Explicit*. Cambridge, MA: Harvard University Press.

Bilgrami, A. (1993). Norms and Meaning, in Ralf Stoecker (ed.) *Reflecting Davidson*, Berlin: de Gruyter.

Blackburn, S. (1984). The Individual Strikes Back. *Synthese*, *58*, 281-301.

Blackburn, S. (2005). Reply: Rule-Following and Moral Realism, in S. Holtzman and C.M. Leich (eds.). *Wittgenstein: To Follow a Rule*. Abingdon: Routledge.

Blank, A. (2007). Wittgenstein on Expectation, Action, and Internal Relations, 1930–1932, *Inquiry*, *50* (3), 270-287.

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

Boghossian, P. (2003). The Normativity of Content. *Philosophical Issues*, *13*, 31-45.

Boghossian, P. (2005a). Is meaning Normative?. In A. Beckerman and C. Nimtz (eds.). *Philosophy – Science – Scientific Philosophy*. Paderborn: Mentis.

Boghossian, P. (2005b). Rules, Meaning and Intention – Discussion. *Philosophical Studies*, *124*, 185-197

Boghossian, P. (2008). Epistemic Rules. *The Journal of Philosophy*, *105* (9), 472-500.

Boghossian, P. (2012) Blind Rule-following, in Annalisa Coliva (ed.). *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*. Oxford: Oxford University Press.

Brandom, R. (1998). *Making it Explicit*. Cambridge, MA: Harvard University Press.

Brewer, B. (2011). *Perception and its Objects*. Oxford: Oxford University Press.

Bridges, J. (2014). Rule-following Skepticism, Properly So Called, in A. Kern and J. Conant (eds.). *Varieties of Skepticism: Essays after Kant, Wittgenstein, and Cavell*. Berlin: de Gruyter.

Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.

Cavell, S. (1999). *The Claim of Reason*. New York and Oxford: Oxford University Press.

Cavell, S. (2013). *This New Yet Unapproachable America*. Chicago: The University of Chicago Press.

Conant, J. (2002). The Method of the Tractatus, in Erich H. Reck (ed.). *From Frege to Wittgenstein: Perspectives in Early Analytic Philosophy*. Oxford: Oxford University Press.

Conant, J. (2017). Kant's Critique of the Layer-Cake Conception of Human Mindedness in the B Deduction, in James O'Shea (ed.). *Kant's Critique of Pure Reason: a Critical Guide*. Oxford: Oxford University Press.

Connelly, J. (2012). Meaning is Normative: A Response to Hattiangadi. *Acta Analytica*, *27*, 55-71.

Crane, T. (2010). Wittgenstein on Intentionality. *The Harvard Review of Philosophy*, *17*, 88-104.

Crane, T. (2014). Wittgenstein and Intentionality, in *Aspects of Psychologism*. Cambridge, MA: Harvard University Press.

Descartes, R. (1960) [1637]. *Discourse on Method and Meditations*. New York & London: Macmillan.

Descartes, R. (1947) [1637]. *Discours de la méthode*. Paris: Librairie Philosophique J. Vrin.

Descartes, R. (1990) [1647]. *Méditations métaphysiques*. Duc de Luynes (trans.). Paris : Livre de Poche.

Diamond, C. (1988). Throwing Away the Ladder. *Philosophy*, *63*, 5-27.

Dreyfus, H. (2005). Overcoming the Myth of the Mental: How Philosophers Can Profit from the Phenomenology of Everyday Expertise. *Proceedings and Addresses of the American Philosophical Association*, *79* (2), 47-65.

Dreyfus, H. (2006). Merleau-Ponty and Recent Cognitive Sciences, in Taylor Carman and Mark B. N. Hansen (eds.) *The Cambridge Companion to Merleau-Ponty*. Cambridge: Cambridge University Press.

Dreyfus, H. (2007). Reply to Romdenh-Romluc, in Thomas Baldwin (ed.) *Reading Merleau-Ponty: On Phenomenology of Perception*. Abingdon and New York: Routledge.

Dreyfus, H. (2017). Could Anything Be More Intelligible than Everyday Intelligibility?, in Mark Wrathall (ed.). *Background Practices: Essays on the Understanding of Being*. Oxford: Oxford University Press.

Dreyfus, H. & Taylor, C. (2015). *Retrieving Realism*. Cambridge, MA: Harvard University Press.

Garland, E. C. & McGregor, P. K. (2020). Cultural Transmission, Evolution, and Revolution in Vocal Displays: Insights From Bird and Whale Song. *Frontiers in Psychology*, *11*, 544929.

Ginsborg, H. (2010). Review: Oughts and Thoughts: Rule-Following and the Normativity of Content by Anandi Hattiangadi. *Mind*, *119* (476), 1176-1186.

Ginsborg, H. (2011a). Primitive Normativity and Skepticism about Rules. *The Journal of Philosophy*, *108* (5), 227-254.

Ginsborg, H. (2011b). Inside and Outside Language: Stroud's Nonreductionism about Meaning, in Jason Bridges, Niko Kolodny, and Wai-hung Wond (eds.). *The Possibility of Philosophical Understanding: Reflections on the Thought of Barry Stroud*. Oxford: Oxford University Press.

Ginsborg, H. (2012). Meaning, Understanding and Normativity. *Proceedings of the Aristotelian Society Supplementary Volume*, *86*, 127-146.

Ginsborg, H. (2020). Wittgenstein on Going On. *Canadian Journal of Philosophy*, *50* (1), 1-17.

Ginsborg, H. (2021). Going on as one ought: Kripke and Wittgenstein on the normativity of meaning. *Mind & Language*, 1-17.

Glendinning, S. (2007a). *In the Name of Phenomenology*. Abingdon and New York: Routledge.

Glendinning, S. (2007b). The genius of man, in Thomas Baldwin (ed.) *Reading Merleau-Ponty: On Phenomenology of Perception*. Abingdon and New York: Routledge.

Glüer, K. and Pagin, P. (1999). Rules of Meaning and Practical Reasoning. *Synthese*, *117*, 207-227.

Glüer, K. and Wikforss, A. (2009). Against Content Normativity. *Mind*, *118*, 31-69.

Glüer, K. and Wikforss, A. (2015). Meaning Normativism: Against the Simple Argument. *Organon*, *Supplementary Issue*, 63-73.

Goldfarb, W. (1983). I want you to bring me a slab: Remarks on the opening sections of the *Philosophical Investigations*. *Synthese*, *56*, 265-282.

Hacker, P.M.S. (1996). Intentionality: the harmony between language and reality, in *Wittgenstein: Mind and Will. Part 1: Essay*. Oxford: Blackwell.

Hacker, P.M.S. (2013a). Crane on Intentionality. *The Harvard Review of Philosophy*, *19,* 46-58.

Hacker, P.M.S. (2013b). Intentionality and the Harmony between Thought and Reality, in *Wittgenstein: Comparison and Context*. Oxford: Oxford University Press.

Haddock, A. (2012). Meaning, Justification, and 'Primitive Normativity'. *Proceedings of the Aristotelian Society Supplementary Volume*, *86*, 147-174.

Hanks, P. W. (2007). The Content-Force Distinction. *Philosophical Studies*, *134*, 141-164.

Hattiangadi, A. (2003). Makint it Implicit: Brandom on Rule Following. *Philosophy and Phenomenological Research*, *56* (2), 419-431.

Hattiangadi, A. (2006). Is Meaning Normative?. *Mind & Language*, *21* (2), 220-240.

Hattiangadi, A. (2007). *Oughts and Thoughts: Rule Following and the Normativity of Content*. Oxford: Oxford University Press.

Heil, J. (2009). Relations, in Robin Le Poidevin et al. (eds.) *The Routledge Companion to Metaphysics*. London and New York: Routledge.

Heil, J. (2012). *The Universe As We Find It*. Oxford: Oxford University Press.

Kelly, S. (2001). The Non-conceptual Content of Perceptual Experience: Situation Dependence and Fineness of Grain. *Philosophy and Phenomenological Research*, *62* (3), 601-608.

Kelly, S. (2007). What do we see (when we do)?, in Thomas Baldwin (ed.) *Reading Merleau-Ponty: On Phenomenology of Perception*. Abingdon and New York: Routledge.

Kripke, S. (1982). *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Cambridge, MA: Harvard University Press.

McDowell, J. (1978). Are Moral Requirements Hypothetical Imperatives?. *Proceedings of the Aristotelian Society, Supplementary Volumes*, *52*, 13-29.

McDowell, J. (1979). Virtue and Reason. *The Monist*, *62* (3), 331-350.

McDowell, J. (1982). Criteria, Defeasibility & Knowledge. *Proceedings of the British Academy 1982*, 455–79.

McDowell, J. (1983) Aesthetic value, objectivity, and the fabric of the world, in Eva Schaper (ed). *Pleasure, preference and value: Studies in philosophical aesthetics*. Cambridge: Cambridge University Press.

McDowell, J. (1984). Wittgenstein on Following a Rule. *Synthese*, *58*, 325-363.

McDowell, J. (1992). Meaning and Intentionality in Wittgenstein's Later Philosophy. *Midwest Studies In Philosophy*, *17* (1), 40-52

McDowell, J. (1996). *Mind and World*. Cambridge, MA: Harvard University Press.

McDowell, J. (1998). Reply to Commentators. *Philosophy and Phenomenological Research*, *58* (2), 403-431.

McDowell, J. (2000) Non-Cognitivism and Rule-Following, in Alice Crary and Rupert Read (eds). *The New Wittgenstein*. London and New York: Routledge.

McDowell, J. (2009). Wittgensteinian Quietism. *Common Knowledge*, *15* (3), 365-372.

McDowell, J. (2011) Values and Secondary Qualities, in Ted Honderich (ed). *Morality and Objectivity: A Tribute to J. L. Mackie*. London: Routledge.

McDowell, J. (2013a). Are Meaning, Understanding, etc., Definite States?, in *The Engaged Intellect*. Cambridge, MA: Harvard University Press.

McDowell, J. (2013b). How Not to Read *Philosophical Investigations*: Brandom's Wittgenstein, in *The Engaged Intellect*. Cambridge, MA: Harvard University Press.

McDowell, J. (2013c). What Myth?, in *The Engaged Intellect*. Cambridge, MA: Harvard University Press.

McDowell, J. (2013d). Perceptual Experience: Both Relational and Contentful. *European Journal of Philosophy*, *21* (1), 144-157.

McGinn, C. (1984). *Wittgenstein on Meaning*. Oxford: Basil Blackwell.

Merleau-Ponty, M. (2004) [1945]. *Phénoménologie de la perception*. Paris: Gallimard.

Merleau-Ponty, M. (2002) [1945]. *Phenomenology of Perception*. Colin Smith (trans.). London and New York: Routledge.

Merleau-Ponty, M. (2012) [1945]. *Phenomenology of Perception*. Donald. A. Landes (trans.). London and New York: Routledge.

Merleau-Ponty, M. (2013) [1942]. *La structure du comportement*. Paris : Presses Universitaires de France.

Miller, A. (2015). Blind Rule-following and the 'Antinomy of Pure Reason'. *The Philosophical Quarterly*, *65* (260), 396-416.

Miller, A. (2017). Rule-Following, Meaning and Primitive Normativity. *Mind*, *0* (0), 1-26.

Millikan, R. (1990). Truth Rules, Hoverflies, and the Kripke-Wittgenstein Paradox. *The Philosophical Review*, *99* (3), 323-353.

Mulhall, S. (2001). *Inheritance and Originality*. Oxford: Oxford University Press.

Mulhall, S. (2007). *Philosophical Myths of the Fall*. Princeton: Princeton University Press.

Mulhall, S. (2014). Inner Constancy, Outer Variation: Stanley Cavell on Grammar, Criteria, and Rules, in A. Kern and J. Conant (eds.). *Varieties of Skepticism: Essays after Kant, Wittgenstein, and Cavell*. Berlin: de Gruyter.

Peacocke, C. (1998). Nonconceptual content defended. *Philosophy and Phenomenological Research*, *58* (2), 381-88.

Pettit, P. (2002). *Rules, Reasons and Norms*. London: Oxford University Press.

Pettit, P. (2016). Broome on reasoning and rule-following. *Philosophical Studies*, *173*, 3373-3384.

Putnam, H. (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.

Read, R. (2021). *Wittgenstein's Liberatory Philosophy: thinking through his 'Philosophical Investigations'*. London: Routledge.

Romdenh-Romluc, K. (2007) Merleau-Ponty and the power to reckon with the possible, in Thomas Baldwin (ed.) *Reading Merleau-Ponty: On Phenomenology of Perception*. Abingdon and New York: Routledge.

Rosen, G. (1997). Who Makes the Rules Around Here?. *Philosophy and Phenomenological Research*, *57* (1)*,* 163-171.

Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.

Sellars, W. (1997) [1956]. *Empiricism and the Philosophy of Mind*. Cambridge, MA: Harvard University Press.

Verheggen, C. (2015). Towards a New Kind of Semantic Normativity. *International Journal of Philosophical Studies*, *23* (3), 410-424.

Whiten, A., Goodall, J., McGrew, W. C., Nishida, T., Reynoldsk, V., Sugiyama, Y., et al. (1999). Cultures in chimpanzees. *Nature*, *399*, 682–685.

Wikforss, A. (2001). Semantic Normativity. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *102* (2), 203-226.

Wittgenstein, L. (1967). *Zettel*. Berkeley, Los Angeles, and London: University of California Press.

Wittgenstein, L. (1969). *The Blue and Brown Books* (2nd Edition). Oxford: Blackwell.

Wittgenstein, L. (1975). *Philosophical Remarks*. Chicago: The University of Chicago Press.

Wittgenstein, L. (1983) [1956]. *Remarks on the Foundations of Mathematics*. Cambridge, MA and London, England: The MIT Press.

Wittgenstein, L. (1990). *Last Writings on the Philosophy of Psychology, Volume 1*. Oxford: Blackwell.

Wittgenstein, L. (2001) [1961]. *Tractatus Logico-Philosophicus*. D.F. Pears and B.F. McGuiness (trans.). Abingdon and New York: Routledge.

Wittgenstein, L. (2009). [1953] *Philosophical Investigations* (Fourth Edition). Oxford: Wiley-Blackwell.

Wrathall, M. (2006) Motives, Reasons and Causes, in Taylor Carman and Mark B. N. Hansen (eds.) *The Cambridge Companion to Merleau-Ponty*. Cambridge: Cambridge University Press.

Wrathall, M. (2007) The phenomenology of social rules, in Thomas Baldwin (ed.) *Reading Merleau-Ponty: On Phenomenology of Perception*. Abingdon and New York: Routledge.

Wrathall, M. (2017) Introduction, in Mark Wrathall (ed.). *Background Practices: Essays on the Understanding of Being*. Oxford: Oxford University Press.

Wrathall, M. & P. Londen (2019) Anglo-American Existential Phenomenology, in Kelly Becker and Ian D. Thomson (eds.). *The Cambridge History of Philosophy, 1945—2015*. Cambridge: Cambridge University Press.

Wright, C. (1984). Kripke's Account of the Argument Against Private Language, *The Journal of Philosophy*, *81*, 759-777.

Wright, C. (1996) Human Nature?. *European Journal of Philosophy*, *4* (2), 235-254.

Wright, C. (2001). *Rails to Infinity*. Cambridge, MA: Harvard University Press.

Wright, C. (2007). Rule-following without Reasons: Wittgesntein's Quietism and the Constitutive Question. *Ratio*, *20* (4), 481-502.

Wright, C. (2012). Replies Part I: The Rule-Following Considerations and the Normativity of Meaning, in Annalisa Coliva (ed.). *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*. Oxford: Oxford University Press.