

Meaning as a Normative Stance

Arnaud Petit

A thesis submitted in partial fulfillment of the requirements for the
Master's degree in philosophy

Department of Philosophy
Faculty of Arts
University of Ottawa

© Arnaud Petit, Ottawa, Canada, 2017

Table of Contents

Acknowledgments	iii
Abstract	iv
0. Introduction	v
1. Meaning, Understanding and Intelligent Behaviour	1
1.1. Introduction	1
1.2. The Intentional View on Rule-Following	2
1.3. Rule-Following as a Matter of Understanding	7
1.4. Must There be Something Behind the Façade of Behaviour?	10
1.5. Mechanical Versus Intelligent Behaviour	13
1.6. The Normativity of Intelligent Behaviour	20
1.7. Toward a New Account of the Normativity of Meaning	23
2. The Normativity of Meaning: Rethinking the Natural	25
2.1. Introduction	25
2.2. Correctness Conditions	27
2.3. The Fairground Example	31
2.4. Prescriptions	36
2.5. Correct as Appropriate	41
2.6. A New Threat to Naturalism?	50
2.7. Conclusion	53
3. Conclusion	55
3.1. Rule-Following, Conditions of Correct Use and Obligations	55
3.2. Salvaging Normativism	58
3.3. Moving Forward	61
Bibliography	66

Acknowledgements

I would like to sincerely thank my supervisor, Professor Patrice Philie, for the role he played in my philosophical development. I owe many of the insights that guide what follows to discussions we had over the last five years. He made me discover Cavell, pushed me to refine my reading of Wittgenstein and ultimately, offered me philosophical tools to better approach the question of the normativity of meaning. I can safely say that he has made me a better thinker and for that, I am indebted to him.

I would also like to thank Eric Wilkinson, John Atytalla, and Pierre-Yves Rochefort for having read and commented various drafts of the following two articles. My gratitude goes also to Professors Vincent Bergeron and Nigel De Souza with whom I had many stimulating exchanges.

I must specially thank Rabia Abdeddaïm for her unwavering support and her constant encouragements. Without her, I could not have done it.

Finally, I am grateful to the Social Sciences and Humanities Research Council of Canada (SSHRC) and the Fonds de Recherche du Québec en Société et Culture (FRQSC) for their funding.

Abstract

In the past few years, the claim that meaning is normative has grown increasingly suspect and many powerful arguments have been developed against its interpretation in terms of (1) conditions of correct use, (2) prescriptions and (3) rule-following. In the first essay of my thesis, I discuss the precise arguments that have been invoked by Paul Boghossian and by Kathrin Glüer and Asa Wikforss against the latter interpretation. In the second essay, I turn towards the two other interpretations of the normativity of meaning, as they are discussed by Anandi Hattiangadi and by Daniel Whiting. My main strategy in both of these essays is to show that the possibility of following a rule, like the existence of conditions of correct use, depends on our ability — and our willingness — to adopt a normative stance. That is to say, I defend the normativity of meaning by showing (i) that we (as human beings) are not indifferent as to how things are done or how words are used and (ii) that the notions of “rule-following” and of “conditions of correct use” are ways to express that fact. In the first essay, I attempt to clarify what I mean by adopting a normative stance — and following a rule — by linking them with the notions of understanding and of intelligent behaviour. In the second essay, I rather insist that conditions of correct use can be explained in terms of appropriateness or of fittingness with the circumstances. As such, I put myself in a position to criticize the traditional unfolding of that notion and to further articulate my idea of a normative stance.

Introduction

In *Wittgenstein on Rules and Private Language*, Kripke develops a famous argument against the possibility of meaning. The argument goes as follows: suppose that a speaker masters the use of the sign '+', but never had to add numbers greater than 56. What justifies him in answering '125' when asked '57+68=?' and not '5'? Or to put it differently, how can he know that he has previously understood – or meant – *addition* by '+' and not, let's say, *quaddition*¹? Indeed, if all his past uses of the sign '+' are compatible with the fact that he means addition, they are also compatible with the fact that he means quaddition. It thus appears that a speaker can never mean anything by a given sign since what he means by it is always under-determined by his past uses of the sign.

Kripke, of course, discusses the most obvious replies to this argument. Among them he discusses the suggestion that the skeptical conclusion – that meaning is not possible – can be avoided if meaning is conceived as a disposition (i.e. as a predictable tendency to respond in a certain way). Yet, according to Kripke, this response is clearly misguided since:

The dispositionalist gives a descriptive account of the relation [between meaning and use]: if '+' meant addition, then I will answer '125'. But this is not the proper account of the relation, which is normative, not descriptive. The point is not that, if I meant addition by '+', I *will* answer '125', but that, if I intend to accord with my

¹ Kripke defines quaddition (+') as followed: if x and $y < 57$, $x+'y = x+y$. Otherwise (that is to say, if x or $y \geq 57$), $x+'y = 5$ (see Kripke, 1982, 9).

past meaning of ‘+’, I *should* answer ‘125’ [...] The relation of meaning and intention to future action is normative, not descriptive (Kripke, 1982, 37)

This passage alone has probably been as influential as the skeptical problem itself. Indeed, in the years that followed the publication of Kripke’s book, many authors were impressed by normativism (i.e. the idea that the relation between meaning and use is essentially – and irreducibly – normative). As Daniel Whiting puts it: at the time, the idea “looked close to achieving the status of orthodoxy” (Whiting, 2009, 535). It was seen as a basic requirement on any theory of meaning (See for example Blackburn, 1984; 286-287; Boghossian 1989, 513; McDowell, 1984; 336; McGinn, 1984, 147 and 174; Wright, 1989). Yet, it is not the case anymore. Normativism has grown increasingly suspect and many powerful arguments have been developed in recent years against its different versions, most notably by Boghossian and the Stockholm School (see for example Boghossian 2005a and 2008; Glüer and Pagin, 1999; Glüer and Wikforss, 2009 and 2015; Hattiangadi 2006 and 2009; Wikforss, 2001). As a matter of fact, with the notable exception of Hannah Ginsborg and Daniel Whiting (see especially Ginsborg 2010, 2011 and 2012 and Whiting, 2006, 2009 and 2016), few authors now defend the normativity of meaning.

The two articles that follow are situated within that discussion. I agree with Boghossian and the Stockholm School that normativism, as it has been interpreted by the first commentators of Kripke’s book, faces important – if not insurmountable – difficulties. Ginsborg’s and Whiting’s respective positions are not entirely satisfying, but they have the merit of showing how normativism can potentially be salvaged. The goal of my two articles is to address and answer the recent challenges that have been raised against normativism. To do so, I rely, of course, on Ginsborg’s and Whiting’s works, but I also explore Wittgenstein’s later writings and Cavell’s discussion of it. Ultimately, I attempt to show

that an original account of the normativity of meaning can emerge from these considerations.

My thesis is divided as follows: in a first article, I show that the notion of following a rule, adequately construed, can be used to defend the normativity of meaning. To do so, I address the arguments that have been raised against this very strategy by Boghossian (2008 and 2009), Glüer and Pagin (1999) and Glüer and Wikforss (2009). I also show how my own account of following a rule can be distinguished from the one proposed by Ginsborg (2011 and 2012). In a second article, I argue that Hattiangadi (2006) – and Glüer and Wikforss (2009 and 2015) – are mistaken when they claim that the normativity of meaning cannot be accounted for in terms of correctness conditions or in terms of prescriptions. I suggest that their position is flawed because it relies on contentious assumptions about language and normativity. I also discuss the question of whether the normativity of meaning necessarily poses a threat to naturalism – as it is assumed by many. Finally, in a conclusion following these two articles, I briefly discuss how I aim to further develop my own account of the normativity of meaning.

Meaning, Understanding and Intelligent Behaviour

Introduction

In *Wittgenstein on Rules and Private Language*, Kripke takes a stand against dispositional accounts of meaning. He insists that meaning – or rather the relation between meaning and use – is normative and that it cannot, as such, be captured by a disposition. That is to say, for Kripke, any attempt to reduce the meaningful employment of an expression to a regularity in behaviour – that is, to a predictable tendency to respond in a certain way – is ultimately misguided since it overlooks an essential feature of the phenomena (i.e. its normative character). Arguably, any other reductively naturalistic account of meaning would face similar difficulties and thus would show itself to be equally misguided. Given such a radical conclusion, it is not surprising that Kripke’s argument – and the insight on which it hinges – has been widely discussed. Commentators were, at first, largely sympathetic toward Kripke’s antireductionist strategy, most of them seeing the normativity of meaning as self-evident. But it is a different story today. The normativity of meaning has grown increasingly suspect and many arguments have been developed 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.²

Of course, Kripke’s insight can be – and, as a matter of fact, has been – interpreted in various ways. For example, as some have suggested, the normativity of meaning could

²The most prominent figures of anti-normativism are probably Boghossian and authors I have described in the introduction as the Stockholm School, namely Glüer, Hattiangadi, Pagin and Wikforss. (See, for

simply consist in the fact that meaningful expressions have correct and incorrect uses (see Blackburn 1984 and Boghossian 1989). But it could equally amount to the fact that the meaning of an expression implies obligations about how the expression ought to be used. Finally, a case could be made for understanding Kripke's insight as implying that, in using an expression meaningfully, one is following a rule. Yet, none of these ways to interpret the normativity of meaning has escaped criticism.

In the present essay, I want to discuss the latter interpretation of Kripke's insight. I will take as my starting point a number of concerns that have been recently raised about the possibility of making sense of the normativity of meaning in terms of rule-following. I will show that Hannah Ginsborg's account of rule-following – and of language use – successfully answers these concerns (see Ginsborg 2011 and 2012), but faces, in turn, other important difficulties. My position is that Ginsborg is right in suggesting that rule-following requires *understanding*, but wrong to assume that this understanding amounts to something 'going on' in the mind. I will argue that we can preserve Ginsborg's insight about understanding if we conceive of it as manifesting itself in behaviour. To conclude, I will suggest that a new account of the normativity of meaning can emerge from these considerations.

The Intentional View on Rule-Following

It is uncontroversial, I think, that following a rule cannot consist merely in acting in accordance with a rule if it is to be normative. 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

example Boghossian (2003), Glüer and Pagin (1999), Glüer and Wikforss (2009), Hattiangadi (2006, 2007), Wikforss (2001)).

being familiar with the rule. In jotting numbers down on a piece of paper, 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.³ Likewise, while playing on an empty board, a child might move a rook according to the rules of chess even though he has 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 he is guided by the rule, but because he has been submitted to some kind of brute conditioning. In all these scenarios, the agent does act in accordance with the rule, but he does so without any consideration about what he ought to do, about whether or not what he does is correct. He is ultimately not responsive (or so it seems) to a norm, even if he happens to satisfy its demands. Of course, in such scenarios, the behaviour of the agent could display certain regularities, but even these regularities would not be sufficient to establish that he is following a rule. The point is that even if the agent were to conform again and again to a rule – even if he came to exhibit a disposition to do so –, we could still entertain doubt as to whether or not he is following the rule. After all, in the scenarios in which the agent is submitted to brute conditioning, it is not clear how he would be different from an automaton, or, if he was to speak, from a parrot.

What, then, is needed in order to say of the agent that he is doing something essentially normative? What would give us the assurance that he is genuinely following a rule? A common answer is that the agent must see the rule in question as a *reason* for acting as he does. The rule must not only explain his behaviour, it must also be part of his motivation. That is what Glüer and Pagin have in mind when they suggest that following a

³ The example is from Bridges (2014).

rule *R* consists in invoking *R* in contexts of practical reasoning (see Glüer and Pagin, 1999). A rule *R* would not be followed by an agent if it were not internal to his resolution to act in one way rather than another, if it did not guide his actions. 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, in this line of thought, there is, we could say, an intentional requirement on rule-following that distinguishes it from mere conformity to a rule.⁴

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 now allows us to substantiate 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, he is not following the rules of chess, it is because he does not, at any point, *intend* to abide by these rules, being, by stipulation, ignorant of them. Second, it seemingly supports the idea that following a rule is a distinctive human ability, thus reinforcing the long-standing intuition that human beings are fundamentally different from automata or parrots. It suggests that following a rule is a matter of navigating in a 'space of reasons' and that human beings have privileged access to such a space. The idea is that it is only within it

⁴ This requirement is assumed, I think, by a wide range of authors. It has been recently invoked by Boghossian (2005b and 2008), Glüer and Wikforss (2009) and Pettit (2002).

that intentions can be formed and rules invoked as reasons. Since automata, parrots and the like do not navigate in such a space, it makes no sense to say that they follow rules.

Yet, despite these merits, such a line of thought faces important difficulties when we further assume that meaning is itself a matter of rule-following. Indeed, as some authors have pointed out, it is difficult to see how the intentional requirement could then be satisfied (see most notably Boghossian 2008 and Glüer and Wikforss 2009). The general idea is that in trying to satisfy the requirement, we would inevitably embark on some kind of endless regress. For Boghossian, it is rather straightforward: forming an intention to conform to a rule already requires the ability to speak – or at the very least, to think – meaningfully. Yet, under our assumptions, speaking meaningfully is a matter of following a rule, which requires, in turn, an intention to conform to the rule, etc.

Glüer and Wikforss approach the question from a somewhat different angle, but their conclusion is the same. Their argument goes as follows: a rule *R* can only motivate an agent's verbal behaviour, if, from that very rule, the agent can infer what to say. Yet, an agent's commitment to conform to *R* – what Glüer and Wikforss call his pro-attitude toward *R* – is clearly not sufficient to tell him what to say. The agent also needs to know what verbal expression would be in accordance with *R* in the situation. Thus, a rule *R* can only motivate an agent – and so, can only be genuinely followed – if it enters into the following kind of practical reasoning:

(P₁) I intend to conform to *R*

(P₂) To say 'p' is in accordance with *R*

(C) I intend to say 'p'.

The problem is that, if rules are constitutive of meaning, then the meaning of (P₂) would depend upon further rules, whose potency would, in turn, depend upon the recognition of what would be in accordance with them, etc. To accept both constraints on rule-following (i.e. meaning constitutivity and the intentional requirement) would lead to an endless regress.⁵ Thus, for Glüer and Wikforss, the only way to defend the claim that meaning is a matter of rule-following would be to abandon the intentional requirement – and so to recognize that a rule could be followed without any commitment to it. However, 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, 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, making sense of the normativity of meaning in terms of rule-following appears to be condemned to failure.

The passage by Glüer and Wikforss just quoted reveals, however, that such a conclusion relies largely on the assumption that the intentional requirement is the only way to make sense of rule-following. With Boghossian, they claim that such an assumption is “overwhelmingly intuitive”, but their certainty on the matter appears to be somewhat groundless.⁶ As I will argue in the next section, Ginsborg’s defence of the normativity of

⁵ Glüer and Wikforss call it the “regress of motivation” (Glüer and Wikforss, 2009, 56).

⁶ Now, it should be noted that Boghossian’s position on that matter is not as clear nor as fixed as Glüer and Wikforss. In his 2005 discussion of Pettit’s *Rules, Reasons and Norms* and of rule-following, he recognised that the intentional requirement is incompatible with the meaning constitutivity of rule-following. But he seems willing to abandon the requirement (which, as he himself points out, was not the case in his 1989 essay on the question). In his 2008 essay, he suggests rather that we might ultimately be forced to abandon

meaning suggests a rather different way to substantiate our intuitions about rule-following – and thus gives hope that the dilemmas exposed by Boghossian and Glüer and Wikforss can be avoided.

Rule-Following as a Matter of Understanding

Ginsborg suggests that what distinguishes following a rule from merely acting in accordance with a rule is that the former presupposes understanding 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 *R*, 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 a sense 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 Wittgenstein’s famous language-game of the builder and his assistant⁷ 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 such a disposition would not be sufficient to establish that the assistant understands what is being asked of him. 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 intentional requirement independently of any consideration about meaning. Still, in this essay, he seems to think that in doing so, we might end up an empty notion of rule-following.

⁷ See Wittgenstein, *PI*, § 2.

the order “Slab!”) is not sufficient to establish that an agent is following a rule since the behaviour of an automaton – clearly lacking the ability to abide by a norm – can still exhibit this kind of regularity. But, for Ginsborg, it would be a different story if, in addition to predictively 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 feeling of appropriateness accompanies the assistant’s actions – that we can legitimately talk of rule-following.

Now Ginsborg insists that this sense 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 see his response to the order “Slab!” as appropriate – or correct –, even though he does not recognize, 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 memorized 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. Many of our everyday activities are indeed governed by rules, but, for Ginsborg, this observation should not obscure the fact that it is ultimately possible to see one’s own action as appropriate *simpliciter* – and so, to follow a rule without any reference to an “antecedently recognized rule” (Ginsborg, 2011, 233).

For Ginsborg, these considerations on rule-following also apply to language use (which can be thought of as a particular case of rule-following). 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 “red”. There is thus no sense in saying that such rules guide their uses of “red”. But for Ginsborg, they are still employing the word “red” meaningfully if, whenever they do so, they have the feeling of employing it appropriately. Nothing more is needed to distinguish competent speakers from the child who, while babbling, produces the sound “red”, or from the parrot trained to utter the sound “red” whenever it sees something red. If it has been properly trained, the parrot will, time and again, conform to certain rules about how to use “red”, but, like the child, it will ultimately lack a sense of going on appropriately. And for Ginsborg, this is sufficient to deny both the parrot and the child the status of speakers – and so of rule-followers (as far as the use of “red” is concerned, of course).

Hence, for Ginsborg, rule-following – including language use – does not ultimately depend upon the acceptance – 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).⁸ 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 or, more elaborately, in terms of practical reasoning. As

⁸ The sense of appropriateness proposed by Ginsborg is also ‘primitive’ in the sense that it is constitutive of meaning without being essentially semantic. That is to say, for Ginsborg, the same kind of normativity that is essential to language can also be constitutive of other aspects of human life. (I will return to this point later).

a matter of fact, she even explicitly criticizes recent attempts to do so (see Ginsborg, 2011, 233 and Ginsborg, 2012, 137, respectively). In insisting that rule-following is a matter of understanding – and that it involves a primitive normativity –, she is advocating for an alternative to the intentional view and its central requirement. She is proposing a new way to make sense of the idea that following a rule – or employing an expression meaningfully – is an essentially normative activity. 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 – it will be the focus of the next section. But we can already see how, at the very least, it can avoid the regresses identified by Boghossian and by Glüer and Wikforss.

Indeed, as I have already suggested, we only face these intractable difficulties if we assume both that rule-following is solely explainable in terms of intentions and that it is constitutive of meaning generally. By rejecting the first assumption, Ginsborg opens up the possibility of defending the second without embarking on such regresses. And this is, I will argue, one of the great merits of her account of rule-following. It shows us a way to think of the meaningful employment of linguistic expressions as cases of rule-following without giving flank to what Boghossian and Glüer and Wikforss take to be decisive objections. It shows us that, by freeing ourselves from the intentional requirement, we can preserve the idea that normativity pervades meaning.

Must There be Something Behind the Façade of Behaviour?

However, Ginsborg's account of rule-following also raises certain worries. She might very well be right in suggesting that the meaningful employment of linguistic expressions – and rule-following activities more generally – require some kind of understanding. And

she offers a convincing case in favour of such a requirement. The main issue, in my opinion, concerns rather her insistence that this understanding ultimately amounts to a sense of appropriateness. Indeed, in claiming such, Ginsborg seems to commit herself to the idea that it is something essentially private that distinguishes following a rule from merely conforming to a rule. She seems to be forced to accept that the distinction hinges on something that is only accessible from the first-person perspective. But surely, this cannot be right. It would imply that, while an agent can know if he is genuinely following a rule – and not, let us say, simply acting in a ‘jack-in-the-box’ manner –, an external observer cannot. It would make rule-following completely opaque to the third-person perspective. The issue is that such a position – and its implications – are clearly at odds with our confidence in many judgments made from the third-person perspective.

Looking at a parrot or an automaton, we are more than willing to say that it lacks understanding, even though we have no access to what takes place within it. We are only external observers, but it does not appear to be an obstacle to our ability to know that it is not really following a rule – in most cases at least. It seems ultimately irrelevant whether or not the parrot – or the automaton – senses that it is acting appropriately, which suggests that our judgments on the matter actually lie on other grounds. Likewise, there are many cases in which we feel that an agent’s behaviour is best described in terms of reflexes or automatisms. For example, take a factory worker who is acting out of habit, staring blankly at the production line and repeating again and again the exact same gestures in a mechanical way. Because he is driven by the repetitive nature of his task, we can easily imagine him failing to notice that the pieces he assembles are defective or that the production line has been momentarily interrupted – both of which would have the effect of rendering his work

useless. Take also a shop clerk who, having been instructed to greet every entering client, comes gradually to repeat the same phrase, like a broken record. It is not only that she lacks enthusiasm. In the end, she also fails to genuinely engage with the people surrounding her, to take interest in the effect of her words, on the responses they might prompt, etc. As such, we can easily imagine her greeting at one point not only new clients, but also clients she has already greeted or who are leaving, her colleagues passing by or even, comically, a dummy behind her. If we were to witness such scenes, we would feel justified in saying, I think, that both the worker and the clerk are acting in a ‘jack-in-a-box’ manner.⁹ Here again, we would not be in a position to know whether or not they have the feeling of conducting themselves appropriately. And still, we would be confident that they lack understanding.

Of course, I am not suggesting that there is no asymmetry between the first-person and the third-person perspectives. It is, I think, a platitude that an agent does not relate to his own actions as an external observer would. I am not suggesting either that there is no doubtful case, no case where the third-person perspective reveals itself to be fallible. I am only noting that, in many cases, no doubt seems to arise. The presence or absence of understanding simply appears obvious, even to an external observer. But, because she equates understanding with something essentially private, Ginsborg cannot accept this. She has no other choice but to claim that we are mistaken in these latter cases, that our confidence is ultimately unfounded. In the end, her account of rule-following is irreconcilable with our third-person judgments about the parrot and the automaton, the worker and the clerk, etc. There is thus a price to pay for accepting it. Of course, we might

⁹ To be clear, we would not be denying that the worker and the clerk are capable of understanding. We would simply acknowledge that, in these precise circumstances, understanding cannot be ascribed.

be willing to pay such a price in order to preserve the understanding requirement – and to make sense of rule-following –, but I do not think we have to.¹⁰ Indeed, it seems to me that an alternative is available to us, such that we can still make sense of the understanding requirement without making it depend on a feeling. If I am right, invoking something in the agent’s mind is not the only way to make intelligible the distinction between acting with understanding and acting without understanding. My suggestion, in a nutshell, is that rule-following does not ultimately require more than acting intelligently – as opposed to mechanically. Understanding does not have to be something hidden that accompanies behaviour. It can be seen as fully manifesting itself in behaviour. Wittgenstein suggests something similar, I think, in a few remarks of *Zettel* in which he comes back to the language-game of the builder and the assistant.¹¹

Mechanical Versus Intelligent Behaviour

Now, there is something peculiar – almost paradoxical – about Wittgenstein’s initial description of the language game. On the one hand, we are inclined to think that it can teach us something essential about understanding.¹² On the other hand, it always appears to fall short of delivering that very teaching, as if too many things were left unsaid. We are prone to look at the builder and the assistant as lacking something resembling understanding, but at the same time, it seems that Wittgenstein’s description is not, in itself, sufficient to warrant such a conclusion. It explains, I think, why Ginsborg (and many

¹⁰ Ginsborg, for example, would be willing, I think, to give up our judgements about the parrot, etc. in order to preserve her account of rule-following. After all, she appears to be aware of the implications of her account (see Ginsborg, 2012, p. 144), but does not seem to be bothered by them.

¹¹ See Wittgenstein, *Z*, §§ 98-108.

¹² See for example Rhees 1959, Goldfarb 1983 and Cavell 2013.

others) feel as if they need to fill in the missing details. As we have seen, for Ginsborg, something ultimately needs to be said about what is going on in the assistant's mind. We need to know if, while responding to the queries of the builder, he has the feeling of going on appropriately. It is only then that we can decide whether or not he is acting with understanding.

But this is not the only way to fill in the details. As Wittgenstein's remarks in *Zettel* seem to suggest, it is possible to shed light on the question of understanding simply by situating the assistant and the builder in a broader context.¹³ What is missing in Wittgenstein's initial description might not be a window onto the assistant's mind, but rather a picture of the surroundings. For example, if we were to suppose that the builder and the assistant are primitive men that lack the ability – or the willingness – to do anything more than shout orders and bring stones, we might be reluctant to credit them with understanding. However, we would not have such doubts if we were to assume that they are very much like us, but that they are working on a construction site full of noise and bustling with activity.¹⁴ Why so? The beginning of an answer might be to note that, in the second case, we can easily imagine the builder and the assistant getting around confusions and misunderstandings. We can imagine them agreeing on better ways to get the work done, attempting to solve problems as they appear, improvising when necessary, etc.¹⁵ When we assume that they are like us, the builder and the assistant appear to be able to act intelligently, despite the inherent limits of their language-game. Thus, if the primitive men

¹³ In the remarks 98-108, Wittgenstein is specifically concerned with the idea that speaking meaningfully implies thinking. He does not explicitly address the question of the relation between understanding and rule-following, but I think the two questions are akin.

¹⁴ Both examples are from Cavell (see Cavell, 2013, 62-64). Goldfarb presents similar examples, but he mentions in a note that they were first suggested to him by Cavell (see Goldfarb, 1983, 269-270).

¹⁵ Cavell and Mulhall suggest similar things (see Cavell, 2013, 63 and Mulhall, 2007, 110-111)

are condemned to carry out their tasks in a mechanical way – as we have suggested –, it is not because they lack words, but because their very nature does not afford them any other alternative.

Wittgenstein is proposing something similar when he claims that there is an important distinction to be made “between creatures that can learn to do work, even complicated work, in a mechanical way, and [creatures] that make trials and comparisons as they work.” Wittgenstein is not clear as to what “trials and comparisons” is supposed to encompass, but he insists that any example would have to be taken from “our life or from a life that is like ours” (Z, § 103). It suggests, I think, that *we* are ultimately the metre stick of what I have called ‘intelligent behaviour’. It does not mean, of course, that other creatures cannot exhibit such behaviour or that we cannot come to behave merely mechanically. It simply makes it clear that intelligent behaviour can be explained in terms of what we do – or are able to do. Now, if we were playing the role of the assistant and we received the order “Slab!” we would be able to:

- correct ourselves upon realizing that we have grabbed a beam and not a slab despite having been asked to bring a slab;
- correct others upon noticing that they have made a similar mistake;
- challenge the calls of the builder if we feel that it is not a slab that is needed at that point but a pillar, etc.;
- identify that a slab is defective and so unfit to be used by the builder;
- recognize a stone as a slab despite minor – and essentially aesthetic – defects that make the stone somewhat dissimilar to the other slabs

And *that* is what responding intelligently to such an order amounts to. Of course, we might occasionally fail to do such things. But if it happened again and again – that is if we systematically failed to take notice of the changing circumstances and to adapt our

responses to the order “Slab!” accordingly –, it would become clear, I think, that we are not acting intelligently anymore, but merely mechanically.

Thus conceived, the distinction between mechanical and intelligent behaviour is particularly well-suited to account for many intuitions we have about rule-following – and language use. It puts us in a position to say, for example, that the automaton and the parrot simply lack the relevant ability to act intelligently – and so to exhibit understanding. Likewise, it allows us to acknowledge that, like the factory worker and the store clerk, we might find ourselves momentarily confined to mechanical gestures and lifeless utterances. (We could say that it allows us to see ourselves as always in danger of becoming a parody of ourselves)¹⁶. Yet, it still makes it clear that, in most cases, we are not, as a matter of fact, behaving in a ‘jack-in-the-box’ way.

At this point, some might feel that I am making too much of the distinction between intelligent and mechanical behaviour. I am in effect suggesting that the distinction is sufficient to make sense of the difference between understanding the orders of the builder and merely being disposed to respond to these orders in a certain way. Yet, it might be argued that what I have sketched as intelligent behaviour is nothing more than further dispositions. For some, it might seem as if correcting oneself and others, etc., is not fundamentally different from being disposed to bring a slab whenever the builder shouts “Slab!”. The point is that, if intelligent behaviour can ultimately be broken down into dispositions – and so into small pieces of mechanical behaviour – then it becomes difficult to see how it can be an expression of understanding.¹⁷

¹⁶ I think Cavell and Mulhall arrive at a similar conclusion (see Cavell, 2013, 64 and Mulhall, 2007, 111).

¹⁷ Ginsborg, for one, suggests something along these lines (see Ginsborg, 2012, 134, note 17).

Now, there is a certain ambiguity about the claim on which this challenge hinges (i.e. the claim that intelligent behaviour can be accounted for in terms of dispositions). It can be understood, of course, as a prediction about what we will ultimately find out about such patterns of behaviour. But it can equally be seen as the expression of a particular requirement on what constitutes a proper explanation of behaviour generally. That is to say, the claim can be interpreted either as an empirical one (in which case it says something about what intelligent behaviour *really* consists in) or as a methodological one (in which case it merely tells us what an explanation of behaviour must look like to count as an explanation at all).

It seems to me that the challenge is appealing, in part, because this ambiguity goes unnoticed, because we fail to distinguish the methodological requirement that we *should* only posit dispositions to explain behaviour from the empirical claim that intelligent behaviour *does* only consist in dispositions. Contrary to the empirical claim, the requirement possesses an aura of certainty that is characteristic of the *a priori*. And so, it is not clear how it could be undermined by what we will find out about behaviour – nor conversely how it could be justified by it. Indeed, the requirement implies that if evidences about the nature of behaviour are evidences at all, it is, minimally, because they are framed in terms of dispositions. It would thus be begging the question to justify it by appealing to the fact that, in accounting for behaviour, we always only refer to dispositions. Of course, one can still endorse such a requirement. But then, it should not come as any surprise that one only finds dispositions – neither should it be seen as informative. Something cannot be at the same time an *a priori* demand and a discovery. And this is precisely the problem in conflating the methodological requirement and the empirical claim. It gives the impression

that, in suggesting that intelligent behaviour can be reduced to dispositions, one says something both (i) informative about such patterns of behaviour and (ii) empirically irrefutable.

What happens then when the methodological requirement and the empirical claim are distinguished? To some, it might seem evident that a proper explanation of behaviour should only posit dispositions. It might be thought that behaviour can only become intelligible at that condition, but, in the end, it is not clear why we should accept this. The idea of a disposition is that of a regularity of behaviour, of a predictable tendency to act – or speak – in a certain way. To ascribe a disposition to someone is to posit that, short of interferences, a particular feature of his immediate surroundings will time and again elicit the same response. And, arguably, this fits nicely with certain patterns of behaviour (I am conceding precisely that in introducing the notion of mechanical behaviour). Still, the question remains: why should we assume that there is no other intelligible pattern of behaviour, that no other sequence of gestures and utterances might be of interest to us? It seems to me that part of the motivation in taking such a stand comes from the impression that it is the only way to securely ground what we do and say in nature. Causal relationships are salient in the case of dispositions; they are a lot less so when it is a question of what I have called intelligent behaviour. To avoid putting human life outside the reach of scientific explanation, it can thus be tempting to assume that dispositions are the only *real* patterns of behaviour and that intelligent behaviour is merely an appearance, emerging from their interactions. However, I think such a concession is unnecessary. The notion of intelligent behaviour is not an attempt to think of human life as outside of nature – or as causally unconstrained. It does hint at another way of carving up the weave of human life, but,

ultimately, it does not talk of anything more than gestures and utterances, than bits of behaviour.

Of course, even if that much is conceded, the empirical side of the challenge still needs to be addressed. There might be no *prima facie* reason to suppose, against the notion of intelligent behaviour, that dispositions are the only real patterns of behaviour. But what if we were to find out that, as a matter of fact, every pattern of intelligent behaviour can be reduced to dispositions? Or to put it differently: what if it became clear that the notion of intelligent behaviour is superfluous in accounting for the complexity and nuance of what we do and say? The suggestion does not presuppose that an explanation in purely dispositional terms *has been* found, but simply that one *could be* found. And as such, it is clearly irrefutable. Still I would like to suggest that the burden of proof should be shifted to the dispositionalist: it is he who should have the responsibility of showing the plausibility of his suggestion. The mere prospect of arriving one day at a purely dispositional account of behaviour should not bar outright other attempts to make sense of what we do and say.

At first sight, it appears evident that human life does not consist merely in the kind of repetitions associated with dispositions. Speaking of intelligent behaviour is precisely a way of accounting for the wide range of situations in which we do not seem to simply be doing the same thing as before. The empirical side of the challenge suggests that even such inventiveness emerges from dispositions. If we appear able to escape the kind of closed horizon to which automatons are so clearly confined, it is only because we display more dispositions than them. My response is that, if we are to take this claim seriously, evidence in favour of it must be presented. Short of that, there is no reason to treat intelligent behaviour as an empty notion, to assume that it is condemned to collapse with mechanical

behaviour into innumerable dispositions. And that is why I think the notion is robust enough to make sense of Ginsborg's insight that rule-following – and language use – are a matter of understanding.

The Normativity of Intelligent Behaviour

Thus far, I have argued that the distinction between intelligent and mechanical behaviour is sufficient to make sense of rule-following. Nothing over and above what we do and say need be posited, contrary to what Boghossian, the Stockholm School and even Ginsborg assume. Still, the aim of this article is not merely to defend rule-following against the charge that it is an idle label. It is also to show that, in salvaging rule-following, we can defend the idea that meaning is normative. And so, I will now say a few words about how to make sense of the normativity of intelligent behaviour.

To begin with, the notion of intelligent behaviour encourages us to take seriously the fact that we are capable of inventiveness, of sensibly adapting to changing circumstances.¹⁸ It is, in part, what I was trying to show in imagining ourselves playing the role of the assistant in Wittgenstein's language-game. Of course, we might find ourselves momentarily disconcerted in new circumstances. If the slabs are arranged differently (being aligned vertically rather than stacked horizontally) or if they are chipped or of an intriguing hue, it might take us some time to recognise that they are what the builder is asking for and fit to be used. Nonetheless we would quickly get over those initial hesitations. It would not take us long to realise that a slab remains a slab even when it is disposed vertically rather than horizontally. Likewise, depending on what the builder is doing with the slab, we would

¹⁸ To some extent, it also reminds us that, when we face unanticipated challenges, we are able to devise new ways to reach our goals, to improvise when everything else fails, etc.

have no difficulty in deciding whether or not minor defects are relevant. The point is that we are never completely helpless when we face new circumstances – even when such circumstances show the limits of our habitual ways of proceeding or of the habitual cues on which we rely. We are not at a loss outside the routines we adopt. All this reveals, I think, that intelligent behaviour is - at least partly - a matter of understanding the purpose of the queries that are made, of our response to them, etc. To do what I have sketched, is to understand not only what a stone must look like to be a slab, but also what kind of use it can be put to, what features would render it useless, etc. An automaton might be designed to bring stones of a certain type in response to the order “Slab!” and it might succeed fairly well in normal circumstances. Yet, because it does not understand the purpose of what it is doing, it cannot always cope when there are unanticipated changes – even if these changes are ultimately trivial.¹⁹

This is, I think, one way to grasp the normativity of intelligent behaviour: contrary to the automaton, we see purpose in what we – and others – do and say.²⁰ We might have learned to respond to the orders of the builder by always being presented with slabs oriented in the same way and of the same colour and dimension. However – and unlike the automaton, prisoner of its design –, we can go beyond these particular circumstances. That is to say, we can amend our ways of responding to the builder’s orders so as to better comply with the purpose of these very orders.

Another way to reveal the normativity of intelligent behaviour is to stress that, contrary to the automaton, we are not indifferent to how things are done – or how words

¹⁹ For example, it is not difficult to imagine the automaton failing to fulfill his task simply because the slabs are not oriented in the same way as before.

²⁰ To be clear, the point is not that we act with purpose (although we often do). It is rather that we are responsive to the purpose of actions and utterances.

are used. In many situations, we willingly endorse certain ways of acting or speaking at the expense of others. We adopt a normative stance towards certain bits of behaviour. It is especially clear when we correct ourselves and others, that is, when we recognise something as a mistake and act to rectify it. But it also manifests itself when, facing new or unusual situations, we make a decision as to what should be done – or said. Our previous discussion of what intelligent behaviour amounts to already gave us some examples of what I have in mind. If the builder, while working on the frame of a new structure, were to ask us to bring a slab, we might question his call, sensing that it is rather a pillar that is needed. That is to say, in light of what the builder is doing, we might insist that the order “Slab!” is incorrect. In another situation, we could be momentarily distracted and mistakenly grab a block rather than a slab (despite having agreed to bring a slab). But if we were to realise our mistake on our way to the builder, we would bring back the block to the stone pile and exchange it for a slab. Clearly, such a response to the query of the builder (from the initial mistake to the appropriate measure to correct it) is not merely a series of mechanical gestures. Bringing a slab – or a block – is not just something we would do blindly. It is something we would assess against the broader background of the builder’s work - and this assessment would ultimately manifest itself in the way we deal with misguided orders and improper responses.

My suggestion is thus that our ability to understand the purpose of what we are doing and saying and our willingness to adopt a normative stance toward it are two facets of the normativity of intelligent behaviour. To be sure, Boghossian, the Stockholm School and even Ginsborg would probably dismiss this picture of normativity. Indeed, despite their disagreements on how to account for rule-following, they all share the intuition that

normativity consists in something going on in the mind. For them, normativity is restricted to the inner. It is something that accompanies behaviour, but that never shows itself through behaviour. They posit intentions and feelings as the only way to account for our difference with the parrot and the automaton, because they cannot accept that the difference might already be visible. In doing so, they give voice to the prejudice inherited from modern philosophy that the outer – in which behaviour unfolds – is essentially neutral and, as such, inhospitable to normativity. They reinforce a strict – and problematic – dualism between inner and outer, between mind and behaviour. The notion of intelligent behaviour is ultimately a way to go beyond these oppositions.

Boghossian, the Stockholm School and Ginsborg have a caricatural view of human life. They ignore its many nuances and fail to appreciate the complex patterns that emerge through it. It is thus not surprising that they cannot accept my suggestion: they are blind to the subtle ways in which our behaviour and that of the automaton – or the parrot – diverge. Of course, on a single occasion, we might bring a slab to the builder the same way the automaton would have – or say “red” exactly in the same tone as the parrot. But, if we look long enough, situations will emerge in which it will become clear that our relation to the order “Slab!” or the colour word “red” is unlike that of the automaton and the parrot (the distinction between intelligent and mechanical behaviour stresses precisely that). And it is then that the normativity of our behaviour would become apparent.

Toward a New Account of the Normativity of Meaning

In this essay, I have argued that Ginsborg’s suggestion that there is an understanding requirement – rather than an intentional requirement – on rule-following allows us to avoid

the regresses exposed by Boghossian and by Glüer and Wikforss. It makes it possible again to hold that language use is a matter of rule-following, without being forced to reduce rule-following to an idle label. But, contrary to Ginsborg, I have tried to show that the understanding requirement does not depend upon something in the mind, but rather draws the line between intelligent and mechanical behaviour. A lot more needs to be said, but as I have hinted at in the last two sections, a case can be made for the idea that intelligent behaviour is not only fundamentally distinct from mechanical behaviour – and so apt to substantiate our insights about rule-following – but also genuinely *normative*. According to the intentional view, the normativity of rule-following – and of meaning – seems to arise from our situatedness within a space of reason. For Ginsborg, it rather arises from a primitive sense of appropriateness. Now we have reasons to believe that a third alternative is available. It is clear that we are not indifferent to how things are done or how words are used. But we do not have to presume that this normative stance depends upon something hidden behind the façade of behaviour. The normativity of rule-following can be seen, on the contrary, as always already expressing itself in what we do and say, in the complex weave of human activities.

The Normativity of Meaning: Rethinking the Natural

Introduction

In *Wittgenstein on Rules and Private Language*, Kripke argues against dispositional accounts of meaning, on the ground that such accounts explain the relation between meaning and use in descriptive terms whereas the relation is really normative. As he famously puts it: to say, for example, that one means addition by ‘+’ does not imply that one will answer “125” to “57+68=”. It implies rather that one ought to answer “125” (see Kripke, 1982, 37). Dispositional accounts, it seems, simply do not have the tools to explain a phenomenon fraught with ‘ought’, a phenomenon that is irreducibly normative. It is thus unsurprising that many authors have taken Kripke’s line of reasoning as implying, more generally, that the normativity of meaning poses a threat to naturalism (see for example Boghossian, 2003, 31; Wikforss, 2001, 218; Whiting, 2009, 535). Indeed, it seems rather uncontroversial that a naturalistic account of meaning would deal in ‘is’ and not in ‘ought’. And it seems equally uncontroversial that the latter kind of statements cannot be derived from the former. Thus, if meaning really implies statements about what one ought to do – as Kripke suggests –, then any attempt to explain it in purely naturalistic – and so descriptive – terms is condemned to fail.²¹

Without objecting to that line of reasoning, Hattiangadi nonetheless attempts to show, in a recent paper, that these concerns are uncalled for. She agrees that the normative

²¹ This is, of course, not to say that every descriptive account of meaning is naturalistic. Wikforss and Glüer emphasise this point (see Wikforss, 2001 and Glüer and Wikforss, 2009).

character of meaning would “yield a presumption against naturalism” if ‘normative’ is understood as ‘prescriptive’ – and so as implying judgements about what one ought to do (see Hattiangadi, 2006, 220-221). Yet, she argues that, in fact, meaning is not normative in that sense. She rightly concedes that ‘normative’ could be – and has been – given another sense: unlike Kripke, many authors understood the normativity of meaning as consisting not in prescriptions, but in *conditions of correct use* (see most notably Blackburn 1984 and Boghossian 1989). But she insists that this latter sense of ‘normative’ is easily reconcilable with naturalism. Thus, Hattiangadi’s conclusion rests on two claims: that we cannot make sense of the normativity of meaning in terms of prescriptions and that the normativity of meaning would be innocuous to naturalism if it were cashed out in terms of correctness conditions.

In the present paper, I will critically examine these two claims – and the arguments introduced in support of them. I will attempt to show that they both spring from a misconception about what correctness conditions amount to and that they are, as such, ultimately misguided. Of course, Hattiangadi is not the only one to err on this matter. In fact, the confusion can be traced back to some of the initial proponents of the normativity of meaning. Moreover, it is still perpetuated today by most authors critical of the normativity of meaning (among which are the leading figures of the Stockholm School) as well as by some of its new advocates (e.g. Whiting). I will say a few words about the problematic appeal of this misconception and suggest a way to escape its grip. In parallel, I will attempt to show that a satisfying account of the normativity of meaning can emerge from these considerations. Finally, I will discuss the implication of this new account for naturalism.

Correctness Conditions

In the years that followed the publication of Kripke's book, it was a prevailing opinion that the normativity of meaning resulted from the existence, for meaningful expressions, of conditions of correct use. At the time, Boghossian and Blackburn were the most influential proponent of this position – although they ultimately offered few arguments in its favour. In his initial discussion, Boghossian expresses the thought as followed:

Suppose the expression 'green' means *green*. It follows immediately that the expression 'green' applies *correctly* only to *these* things (the green ones) and not to *those* (the non-greens). The fact that the expression means something implies, that is, a whole set of *normative* truths about my behaviour with that expression: namely, that my use of it is correct in application to certain objects and not in application to others. (Boghossian, 1989, 513)

The general idea is that, in recognising that a sound – or a series of strokes on a piece of paper – is meaningful, one is also recognising that it is bound by conditions of correct use. The expression 'green', is not a mere noise, nor an arbitrary scribble; it has meaning. And this implies that there are cases in which it is correctly used and, consequently, cases in which it is not. Clearly, it is correct to say of grass or of the burgeoning leaves of a tree that they are green, but it is an obvious mistake to say the same thing of a wall of red brick. For Boghossian, the matter is rather straightforward: the expression 'green', given its usual meaning, is only true of green things. That's why it correctly applies to green grass and burgeoning leaves, but not to walls of red brick. Blackburn claims essentially the same thing. After having noted that "there is such a thing as the correct and incorrect application

of [an expression]”, he promptly adds that to say this is “no more than to say that there is truth and falsity” (Blackburn, 1984, 281).

Thus, for Boghossian as for Blackburn, the distinction between the correct and incorrect uses of an expression like ‘green’ is ultimately a way to capture the fact that the expression is true of some things and false of others. And for both, this fact is, in itself, sufficient to give flesh to Kripke’s insight that the relation between meaning and use is essentially – and irreducibly – normative. By ascribing meaning to an expression – and so by stating what the expression is true of –, one is hereby recognising that it can be correctly or incorrectly used. That is, one is implicitly submitting its use to normative evaluation.

At first sight, it seems that the normativity of meaning, thus conceived, would pose a threat to naturalism. Neither Boghossian nor Blackburn explicitly defend this claim, but a case can arguably be made that it follows directly from their recognition that meaningful expressions can be correctly or incorrectly used. The argument would presumably go along these lines: it is evident that ‘correct’ is an essentially normative term and that it cannot, as such, be reconstructed from non-normative (i.e. descriptive) terms. Since it is equally evident that a naturalistic account of meaning would be purely descriptive, one can only conclude that such an account would overlook an essential feature of meaning.

Unsurprisingly, this line of reasoning closely parallels the one often offered in support of the idea that prescriptions – about what one ought to do or say – would resist a naturalistic reduction. Yet, according to Hattiangadi, only the latter is cogent. What goes wrong, then, when one argues that correctness conditions – rather than prescriptions – pose a threat to naturalism? Hattiangadi clearly does not object to the idea that there are correct and incorrect applications of meaningful expressions. As a matter of fact, she even agrees

with Boghossian and Blackburn that it is a platitude concerning meaning (see Hattiangadi, 2006, 222). She does not object either to the assumption that normative terms are irreducible to descriptive terms, nor to the assumption that naturalism depends upon a renouncement of normative terms.²² Thus, her criticism can only concern the idea that, in discussing meaning, ‘correct’ expresses something normative.

These considerations are important to fully appreciate her position in this debate and to dissipate a potential ambiguity about her claim as to the innocuity of correctness conditions for naturalism. Indeed, despite what it might have seemed at first sight, Hattiangadi is not suggesting that Boghossian and Blackburn have offered a proper account of the normativity of meaning, which happens to be compatible with naturalism. Quite the contrary. What she is saying is that if the normativity of meaning is, as Boghossian puts it, “simply a new name for the familiar fact that [...] meaningful expressions possess conditions of correct use” (Boghossian, 1989, 513), then it is nothing more than an idle label. And that is why, according to her, it is compatible with naturalism. Ultimately, her position is that Boghossian and Blackburn are mistaken and that the existence, for meaningful expressions, of conditions of correct use does not reveal, in any way, something normative about meaning.

For Hattiangadi, Boghossian and Blackburn have been careless. They have rightly revealed that “correct is just a catch all phrase for the various semantic relations terms can have to the world” (Hattiangadi, 2006, p. 224). But they have then failed to appreciate the

²² She willingly concedes, for example, that one cannot derive an ‘ought’ from an ‘is’ (see Hattiangadi, 2006, 226). Moreover, she clearly defines naturalism as a refusal, in principle, to consider things like normativity to be “among the building blocks of the universe” (Hattiangadi, 2006, 220). And so, she is forced to accept that terms that express normativity would be superfluous in a naturalistic account – of meaning or of anything else.

implications of that fact. If to say that ‘green’ is correctly used is simply to say that the thing to which it refers is indeed green, that ‘green’ is true of it, then it is not the expression of a normative evaluation. It is merely the conclusion of an act of sorting between things which are green and things which are not. Indeed, for Hattiangadi, Boghossian and Blackburn’s mistake is evident: semantic relations – such as being true of something – can be fully accounted for in descriptive terms. And so, if conditions of correct use are just a way to capture these relations, then there is no reason to assume that they are normative. Furthermore, there is no reason to assume that the fact that ‘green’ correctly applies only to green things would have any direct implication as to how ‘green’ should be used. After all, as Hattiangadi suggests, while one might recognise that ‘green’ is only true of green things, one might not always wish to speak truthfully. In such cases, it is not clear that one would be violating any obligation he might have. The point is that, for Hattiangadi, no normative constraint on what one should say appears to follow from the fact that an expression has correct and incorrect uses. And since in merely stating the conditions under which an expression is correctly used, one is not saying something normative either, it is difficult to see how the existence of such conditions would support the normativity of meaning.²³

Hattiangadi proposes the following analogy to illustrate her argument: conditions of correct use are like the height requirement enforced at a fairground. They are standards and tend, as such, to be expressed in evaluative terms – such as ‘correct’, ‘right’, etc. –, but they are not thereby irreducibly normative. That a child is a meter tall – and thus satisfies the height requirement of the fairground – is clearly a non-normative fact (see Hattiangadi,

²³ A similar argument can be found in Glüer and Wikforss, 2009 and in Glüer and Wikforss, 2015.

2006, 224). We might want to say that the child is the ‘correct’ height or the ‘right’ height for the attractions of the fairground, but it is clear that his meeting the requirement can be accounted for in purely descriptive terms. Furthermore, it would be misleading to say that the child should be at least a meter tall, independently of his desire to get on the attractions. The height requirement is a standard, to be sure, but it only aims at sorting children into two categories, those that are at least one-meter-tall – and so, tall enough to get on the attractions – and those that are not. Ultimately, it does not have any implication as to what the child should do.

The Fairground Example

Recently, Whiting has argued that Hattiangadi is wrong on several accounts in her initial discussion of the fairground example (see Whiting 2007, 2009). Hattiangadi claims both that the height requirement is not, in itself, the expression of a normative evaluation and that no obligation directly follows from it. But, for Whiting, her discussion of the example does not ultimately support these conclusions. He willingly concedes that the particular fact that a child is over one meter tall can be accounted for in purely descriptive terms. Still, he wonders how it could support the claim that the height requirement is not expressing something normative (see Whiting, 2009) Of course, it is *because* the child is over one meter tall that he satisfies the requirement. But the requirement is not, in turn, a claim about the child’s particular height. That is to say, the requirement:

R₁: Children must be over one meter tall to enter the attractions

does not express the same thing as the claim:

I: Isaac is over one meter tall

Thus, the recognition that the latter is a purely descriptive statement does not permit us to conclude anything about the normative status of the former. According to Whiting, Hattiangadi fails to appreciate this distinction in her initial discussion of the example. She conflates the requirement with the particular feature on which the satisfaction of the requirement depends.²⁴

Moreover, according to Whiting, Hattiangadi falls short of showing that a requirement such as R_1 does not have any normative implication. Of course, Hattiangadi is right to note that even if R_1 was enforced at the fairground, it would not follow that children would have a categorical obligation to be over one meter tall. (The point is obvious enough). However, Whiting insists that R_1 would still be action-guiding. Indeed, if R_1 was enforced, then the child's height would clearly have implication as to whether or not he *may* (or should not) enter the attractions (see Whiting, 2007, 135-136). A child might not care. If he has no interest whatsoever in entering the attractions, it will likely not matter to him whether or not he meets the height requirement. But the fact remains that if he were less than one meter tall and tried to enter an attraction, he would be doing what he *should not*. He would be contravening to the requirement and so, would open himself to sanctions and reprimands. In clear contrast, the requirement would allow children of more than one meter tall to enter attractions as they please. It would remain silent as to what these children *should* do at the fairground, but it would make it clear, at the very least, that it would be appropriate for them to enter the attractions, that they *may* decide to do so. For Whiting, these considerations are sufficient to show that, when a requirement such as R_1 is in force, then its satisfaction is a normative matter.

²⁴ Whiting, following Rosen, talks more generally of a failure to distinguish *correctness* itself from the *correct-making feature* (see Whiting, 2009, 539).

Whiting thus turns Hattiangadi's analogy against her. If conditions of correct use are indeed like the height requirement of the fairground example – as Hattiangadi suggests –, then she has not yet shown that they do not express anything normative. Of course, the fact that the grass is green can be accounted for in purely descriptive terms. Nobody denies that. But the claim that 'green' correctly applies to grass might nonetheless be normative, either because the claim is in itself normative or because it has implications as to how 'green' may (or should not) be used. For Whiting, all this shows that a principle such as:

$$C_1: S \text{ means } F \text{ by } t \rightarrow \text{for all } x (S \text{ applies } t \text{ correctly to } x \leftrightarrow x \text{ is } f)^{25}$$

can perfectly well capture the idea that meaning is normative.

Evidently, Hattiangadi does not accept such a conclusion. In a recent paper, she forcefully defends the soundness of her initial argument against Whiting's objections (see Hattiangadi 2009). Since then, other leading figures of the Stockholm School have also joined the debate in an attempt to put the matter to rest (see especially Glüer and Wikforss 2015). I think these responses to Whiting's objections are rather enlightening and so, I will briefly expose them.

Hattiangadi, for example, accuses Whiting of begging the question in claiming that R_1 has normative implications. As she puts it: Whiting already characterises the requirement enforced at the fairground as “an explicitly normative standard, specifying conditions under which a child may ride” Yet, according to her, the requirement “need not specify conditions under which a child may or may not ride” (Hattiangadi, 2009, 56).

Whiting assumes that R_1 is equivalent to something like:

²⁵ where S is a speaker, F the meaning of the term t and f , the correct making feature of the term t .

R₂: Children are permitted to enter the attractions if and only if they are over one meter tall.

And that is why he sees R₁ as action-guiding – and so, as having normative implications. But, for Hattiangadi, this is not the only way to interpret the height requirement. Indeed, a case could be made that R₁ simply states the relative safeness of the attractions for children, given their respective heights. That is to say, R₁ could be interpreted, not as equivalent to R₂, but rather to:

R₃: Attractions are safe for children if and only if they are more than one meter tall.

Now, for Hattiangadi, it is clear that, contrary to R₂, R₃ is not, in itself, action-guiding. After all, from R₃ and the fact that Isaac is more than one meter tall, one cannot *directly* conclude anything about what Isaac may or should do. One can merely conclude that it is safe for him to enter the attractions (see Hattiangadi, 2009, 56). Thus, if the height requirement is really only a safety standard – which, for Hattiangadi, is more than plausible –, then it does not have normative implications as Whiting assumes.

In light of these considerations, Hattiangadi concludes that there is a disanalogy between conditions of correct use (as captured by C₁) and Whiting's interpretation of the height requirement. For her, it is clear that conditions of correct use are not, in themselves, action guiding. That 'green' correctly applies only to green things states what 'green' is true of. But, as it stands, it does not tell us how to use 'green' – nor even minimally how it is permissible to use it. (It would only have such normative implications if we were to further assume that we have an *obligation* to say what is true). As such, it is more akin to R₃ than to R₂. Glüer and Wikforss suggests something similar. They claim that, in discussing the fairground example, Whiting merely shows what would happen *if* conditions

of correct use were construed normatively, while remaining silent as to *why* they should thus be construed. Glüer and Wikforss willingly concede that the height requirement would have normative implications if it were the expression of a norm effectively stating for whom it is permissible to enter the attractions. However, for them, and despite Whiting's claim to the contrary, the fairground example ultimately revolves around an act of sorting that can be done in the absence of such a norm. After all, as they note, children can be sorted by heights – for safety or other reasons – independently of any considerations about what would be permissible for them to do. Likewise, Glüer and Wikforss suggest that the applications of 'green' can be sorted between those that are true and those that are false “without any norms to that effect being in force” (Glüer and Wikforss, 2015, 68).

Furthermore, and in the same vein, Glüer and Wikforss insists that the distinction between a standard (whether the height requirement of the fairground example or the conditions of correct use of 'green') and the particular feature on which its satisfaction depends is not sufficient to support the claim that conditions of correct use are normative. In introducing the distinction, Whiting might have shown that, despite Hattiangadi's arguments, it is still conceptually *possible* to account for conditions of correct use in normative terms. However, for them, he clearly falls short of showing that they *must* be interpreted along these lines (see Glüer and Wikforss, 2015, 72). It is difficult to see how he could even do so. After all, as they point out, Whiting agrees with Hattiangadi that the basic semantic concepts – of truth and reference – are not normative. His only way to defend the claim that semantic correctness is nonetheless normative would thus be to say that these basic concepts are not sufficient to account for all essential features of meaning. Yet, without further arguments, Glüer and Wikforss remain unconvinced by this

suggestion, seeing in it a vain attempt to save an outdated intuition about meaning. As they put it: given the concession that the basic semantic concepts are not themselves normative, “the need for construing the notion of semantic correctness as normative [becomes] ever so much harder to motivate” (Glüer and Wikforss, 2015, 72).

Prescriptions

Thus, Hattiangadi – like Glüer and Wikforss – remains unconvinced by Whiting’s objections. I have not said anything yet about what I think of their responses (it will be the focus of the next section), but, at this point, it should be clear that they all hinge on the same assumption: that ‘correct’ is just a placeholder for ‘is true of’ or ‘refers to’ – or, equivalently, that semantic correctness is fully captured by the semantic concepts of truth and reference. One might get the impression that neither side is really listening to the other in this debate. (After all, this assumption was already invoked by Hattiangadi in her criticism of Boghossian’s and Blackburn’s account of the normativity of meaning). However, I do not think this would be a fair assessment, if only because, as we have seen, this debate has forced Hattiangadi to clarify her position – and, ultimately, to find a way to render it more convincing. It has given her the opportunity, for example, to dissipate certain ambiguities about her initial discussion of the fairground example. Her response to Whiting’s objections – like the response of Glüer and Wikforss – have also shed some light on the reasons behind her reluctance to construe conditions of correct use normatively. Finally, it has also made it clear why she is so confident when she claims that meaning would only be normative in a sense that would threaten naturalism if it implied prescriptions – or at least permissions.

Indeed, if conditions of correct use can be construed non-normatively, then, clearly, one cannot simply note that ‘green’ correctly applies only to green things to challenge naturalism. For Hattiangadi, one must also show that it follows from the fact that I mean green by ‘green’ that I ought to apply ‘green’ correctly. Now, in her first paper, she examines various principles that could potentially capture this insight and attempts to determine if one of them is in any way plausible. She starts with the following principle:

$P_1: S \text{ means } F \text{ by } t \rightarrow \text{for all } x (S \text{ ought to (apply } t \text{ to } x) \leftrightarrow x \text{ is } f)$

According to P_1 , if, for example, a speaker means green by ‘green’, then he ought to ensure that he applies ‘green’ to something if and only if it is indeed green. The problem is that such a principle is clearly too demanding. One only need to break down the biconditional to see it. Indeed, one gets:

(1) for all $x (S \text{ ought to (apply } t \text{ to } x) \rightarrow x \text{ is } f)$

(2) for all $x (x \text{ is } f \rightarrow S \text{ ought to (apply } t \text{ to } x)$

The latter conditional implies that one ought to apply ‘green’ to *every* green thing – which is, of course, impossible. After all, some green things will always remain unknown to the speaker (whether because they are lost at the bottom of the sea or on a distant planet out of his reach). Since ‘ought’ implies ‘can’, it is clear that (2) is false and so that P_1 must be rejected (see Hattiangadi, 2006, 226-227).²⁶

In her first paper, Hattiangadi thus suggests the following principle as a more plausible candidate to account for the idea that, in using a term meaningfully, a speaker incurs an obligation to use it correctly:

$P_2: S \text{ means } F \text{ by } t \rightarrow \text{for all } x (S \text{ ought to (apply } t \text{ to } x) \rightarrow x \text{ is } f)$

²⁶ The same argument appears in Glüer and Wikforss 2009

Indeed, contrary to P₁, P₂ does not require of a speaker that he apply ‘green’ to every object to which it would be correct to do so (i.e. to every green object). As such, it can be fully satisfied even if certain green things remain, for ever, unknown to him. Now, in her first paper, Hattiangadi interprets P₂ as requiring of a speaker that he use ‘green’ *only* when it is correct to do so. However, as it turns out, P₂ does not ultimately say that: despite appearances, it is not action-guiding.²⁷ It does not imply anything as to how speakers ought (or ought not) to use ‘green’ or any other terms. Upon taking a closer look at P₂, the point appears obvious enough: if x is indeed f (i.e. if the consequent of the conditional is true), then nothing follows as to what S ought to do (i.e. it is indifferent whether or not the antecedent is true or false). Conversely, if x is not f , then S can still mean F by t without hereby incurring an obligation not to apply t to x . Indeed, as Hattiangadi puts it: if the consequent is false and it is not the case that x is f , it only follows that “you *lack* an obligation to apply [t] to x , not that you are obligated *not* to apply [t] to x ” (see Hattiangadi 2009, 58). Thus, P₂ would remain true even if no obligation of any kind followed from the fact that one means green by ‘green’. As such, it might escape the ‘ought’ implies ‘can’ objection, but it clearly falls short of showing that meaning is normative.

In light of these considerations, a third principle has recently been discussed both by Hattiangadi and by Whiting (see Whiting, 2007 and 2009 and Hattiangadi, 2009):

P₃: S means F by $t \rightarrow$ for all x (S ought not to apply t to $x \leftrightarrow x$ is not f)

P₃ states that, whenever a speaker uses an expression meaningfully, he incurs an obligation not to use the expression incorrectly. The speaker may decide to apply— or refrain from applying— ‘green’ to green things and still mean what he usually does by ‘green’. That is,

²⁷ Hattiangadi herself first pointed out this mistake (see Bykvist and Hattiangadi 2007).

doing so is in full compliance with P_3 , but it is not prescribed by P_3 .²⁸ Still, as Hattiangadi and Whiting both note, P_3 remains action guiding since it effectively prohibits the speaker from applying ‘green’ to things which are not.

Yet, contrary to Whiting, Hattiangadi considers that even such a principle must ultimately be rejected. She justifies her position by claiming that the incorrect use of a term is not always the use that we ought not to make. As she puts it: “Sometimes we ought to lie for the sake of a greater good – such as to protect someone from danger. On other occasions, we ought to use words incorrectly simply because it will have a more powerful effect or because it will make people laugh” (see Hattiangadi, 2006, 227). Hattiangadi proposes the following example to illustrate her point: I might want to teach a child to never touch the oven and, to do so, I might come to tell him that the oven is hot even if it is not. For Hattiangadi it is clear that, in such a case, I would be using ‘hot’ as I ought to, even though I am using ‘hot’ incorrectly (i.e. not truthfully). Not only that, but I would still be meaning what I usually do by ‘hot’. Applying ‘hot’ to something which is not (i.e. a cool oven) would not alter its ordinary meaning for me. However, as Hattiangadi remarks, P_3 “makes it a necessary condition of meaning something by a term that a speaker ought to speak the truth” (Hattiangadi, 2006, 227). As such, it rules out in advance precisely this kind of case – where one incurs an obligation to say something incorrect (i.e. false). For Hattiangadi, this proves that P_3 expresses a requirement that is too demanding to be a purely semantic requirement and so, that P_3 must be rejected. Of course, one might defend the notion that we have an obligation to speak truthfully (it might be thought of as a necessary

²⁸ Whiting, for example notes in the same vein that P_3 is equivalent to “ S means F by $t \rightarrow$ for all x (S may apply t to $x \leftrightarrow x$ is f)” (see Whiting, 2009, 544). For him, P_3 thus captures the idea that meaning is more a matter of permission (and incidentally of proscriptions) than it is a matter of prescription.

condition of communication or as a moral imperative). Hattiangadi's point is merely that such an obligation would not be a semantic one – since its violation would neither alter what we usually mean by our terms nor impede, in any way, our ability to speak meaningfully. And if it is not a semantic obligation, it clearly cannot be invoked, as it is in P₃, to defend the normativity of meaning.²⁹

One could object, as Whiting does for example, that the prescriptions – or rather the permissions and proscriptions – that followed from using an expression meaningfully are not categorical obligations, but merely *prima facie* obligations. Of course, in certain situations, I might have to lie and so, to use an expression incorrectly. But it does not necessarily follow that I do not have a semantic obligation not to do so. Indeed, it might simply be that, in that case, my obligation to not use terms incorrectly is trumped by other – more important – obligations (see Hattiangadi, 2006, 231-232). The idea is that, in certain cases, semantic obligations might reveal themselves to be incompatible with other kinds of obligations (whether moral, pragmatic, etc.). One of the conflicting obligations would thus need to be momentarily abandoned in favour of the other. And given the importance of many of these other obligations in the conduct of our life, it is not unreasonable to think that, in such cases, they would often have precedence over purely semantic ones. But that does not mean that we have no semantic obligation, nor that such obligations are never enforced. Indeed, turning back to Hattiangadi's example, I might have an obligation to put children out of harm's way and the only way to do so might be to incorrectly apply 'hot' to a cold oven. All this, however, is perfectly compatible with the fact that I have a *prima facie* obligation not to apply 'hot' to things that are not.

²⁹ A similar argument can be found in Boghossian's more recent writings on the normativity of meaning (see Boghossian, 2003, 39 and Boghossian 2005a, 207).

To this line of reasoning, Hattiangadi responds the following: by definition, *prima facie* obligations are obligations that can only be overridden by other obligations. Yet, it seems that the alleged semantic obligation to only speak truthfully that P₃ attempts to capture can be overridden “not only by an obligation to lie, but also [...] simply by the desire to lie” (Hattiangadi, 2006, 232). As such, it would be misleading to call it an obligation at all. Knowing that ‘green’ correctly applies only to green things, I could nonetheless apply it to something which is not, on a whim or simply because it amuses me. In such a case, it is clear that I would not be motivated by any obligation (i.e. by doing so, I would not be contributing to a child’s safety, etc.). Yet, for Hattiangadi, it remains that I would still mean what I ordinarily do by ‘green’ - even though I have thus violated my alleged semantic obligation to not speak falsely on the basis of a mere desire to speak falsely. For Hattiangadi, this settles the matter: of course, a naturalistic explanation of meaning would not be possible if meaning implied obligations – or at least permissions and proscriptions –, but the failure of P₁, P₂ and P₃ to account for that insight suggests that no such obligation can be found.

Correct as Appropriate

Hattiangadi’s dismissal of *prima facie* semantic obligations has not convinced Whiting who persists in holding his position. Despite Hattiangadi’s claim to the contrary, he maintains that we do, in fact, incur *prima facie* obligations and so that if we were to speak incorrectly simply because it amuses us, we would be doing something we ought not to (see Whiting, 2009, 550-551). He seems to believe that it is the best way to salvage P₃ – and ultimately to defend the normativity of meaning. I beg to disagree. As I see it, Whiting

is conceding too much to Hattiangadi and, as such, is missing a more decisive objection to her position.

Ultimately, Hattiangadi's whole line of reasoning is problematic because it derives from an inadequate picture of language. Indeed, as we have seen, Hattiangadi insists again and again that, in discussing meaning, 'correct' should be understood as standing for 'refers to', 'is true of' and the like. Yet, it is clear that not all meaningful expressions refer to something or are true of something. Questions and imperatives, wordplays, prayers, salutations can clearly be meaningful even though they lack such semantic properties.³⁰ Hattiangadi's conception of meaning thus appears to be too narrow. Still, as we have established, she is not the only one to err on the matter. This reductive picture of language has also been accepted by the other leading figures of the Stockholm School and even by some of the most prominent advocates of the normativity of meaning (i.e. Blackburn and Boghossian). Even when they are not as explicit as Hattiangadi, the examples of meaningful expressions they tend to invoke remain eloquent. They only talk of terms like 'square', 'green' and 'horse'³¹ - that is, of terms that fit nicely with the idea that meaning is purely a matter of reference. They take them as paradigmatic cases, but in doing so, they fail to appreciate the many other ways in which an expression might be meaningful in – ordinary – language.

The point is not new – it is, after all, a central teaching of Wittgenstein, Austin and the ordinary language philosophers –, but it is nonetheless relevant to our present concern. Indeed, as we have seen, Hattiangadi's whole position rests on the dubious assumption that

³⁰ Whiting has recognised that much recently (see Whiting, 2016, 22, footnote 9). However, he seems to have failed to appreciate its implication for Hattiangadi's arguments.

³¹ see respectively Hattiangadi, 2006, 223; Boghossian, 1989, 513; Wikforss, 2001, 205

the conditions of correct use of an expression are always its conditions of truth. It is what permits her to say, in discussing meaning, that (1) ‘correct’ is not an essentially normative term and that (2) there are cases where we ought to say something incorrect (i.e. not true) because the circumstances demand that we lie, etc. And, as we have seen, it is from (1) and (2) that Hattiangadi concludes both that the existence of correctness conditions poses no threat to naturalism and that meaning does not imply prescriptions (or permissions and proscriptions) as to what one ought or ought not to say.

Yet, these conclusions could be avoided if ‘correct’ were understood more generally, in a way that could account for the referential role often played by terms like ‘square’, ‘green’ and ‘horse’, while at the same time enabling us to make sense of the meaning of expressions such as ‘good morning!’, ‘How have you been?’, ‘thank you!’. One way to do it would be, I think, to insist that in claiming that an expression is correctly – or incorrectly – used, one is stating something about the appropriateness – or inappropriateness – of the expression in the given circumstances.³²

Hence, for example, if it is correct to say ‘good morning!’ to a colleague upon arriving at work, it is because the expression is appropriate given the circumstances, that is, given the time of the day, the social setting, your acquaintance with the person, etc. Conversely, if it is incorrect to yell ‘good morning!’ to every person on the street in the middle of the night, it is because such an expression is not appropriate in these circumstances. The idea of appropriateness fits also nicely with expressions like ‘square’. Imagine a teacher pointing at a square figure and asking his students what shape it is. In these circumstances, ‘square’ would be the correct, that is, appropriate, answer. However,

³² I borrow the concept of ‘circumstances’ from Wittgenstein’s later writings and especially from Cavell’s discussion of these writings. (see Wittgenstein, 2009 and Cavell, 1999).

it would not be if the teacher was instead pointing at a round figure. In this case, the particular figure to which the teacher is pointing is, of course, part of the circumstances. As is the particular context in which the question is asked (i.e. it is asked to children and not to competent adult speakers, in the context of a classroom, maybe at the end of a lesson on basic geometric figures, etc.). Indeed, if the context were different, it might not be as clear what the correct answer would be. For example, following a discussion about how basic geometrical figures are not really found in nature, that they are an idealisation, a theoretical construction, if a teacher were to ask, pointing at a 'square' figure, 'what shape is it?' the appropriate answer might not be 'square', but maybe something like 'it is almost a square, but it is not really one'. The question in both situations is the same. We can even suppose that the teacher is pointing to the same square figure in both cases. Nonetheless, the appropriate answer differs.

If 'correct' is thus understood, then it is not clear why we should accept (2). As we have seen, Hattiangadi insists that, in some circumstances, I might have to lie (and so to use a term 'incorrectly') to ensure a child's safety or to make people laugh. But if what is correct is what is appropriate (i.e. fitting with the circumstances), then it appears that any obligation I might have could, in principle at least, match what it is correct to say. Of course, in applying 'hot' to a cold oven, I would always be saying something false. But, depending on the circumstances, I would not necessarily be saying something incorrect. Indeed, few will deny, I think, that it would be appropriate to use 'hot' like in Hattiangadi's example to teach a child not to touch the oven. It would be especially clear if we were further to assume that the child in question has an irresistible tendency to put his hands on the oven, that he is tall enough to do so (i.e. that the oven is really a potential danger for

his safety), that he has already burnt himself and thus understands the warning not to touch hot things, that nothing else has discouraged him so far, etc. It is not to say that I would hereby incur an obligation to say of the cold oven that it is hot, but simply that it would become permissible to do so. The point ultimately is that, in such circumstances, no competent speaker would seriously disapprove of my use of ‘hot’ or come to believe that I do not mean what we ordinarily do by it. As such, the permission I have to lie in that case is not, as Hattiangadi assumes, a permission to say something incorrect.

The same goes when we use sarcasm or irony, when we invoke metaphors or, more generally, when we transposed terms into new – non-literal – contexts. For example, my fiancée could be asking me again and again “Are you sure you have turned off the oven? Are you sure it is not still hot?” either in an attempt to tease me or simply because she is distracted and repeatedly forgets my answer. If I were to respond, exasperatedly: “Oh you are right, the oven is hot!” while putting my hand on it, even though it is cold, I would be saying something false by applying ‘hot’ to something which is not. But would I hereby be using ‘hot’ in a way that is proscribed by the very meaning of ‘hot’? I think not. Understanding the meaning of a term like ‘hot’ does not consist merely in knowing what ‘hot’ is true of. It is also to appreciate its many purposes in ordinary language – and, as such, the great diversity of circumstances in which it can appropriately be used. If a speaker were incapable of seeing that ‘hot’ can be used to express a white lie or sarcastic exasperation, that it can be projected into new contexts, etc., then he would not be meaning what we do by ‘hot’. This is not to say, of course, that ‘hot’ would be a mere noise for him – or a meaningless series of strokes on a piece of paper. After all, we can imagine such a speaker – for whom life with words is ultimately reduced to literal use – without taking the

very possibility of language away from him. The point is simply that such a speaker would not have the same sensibility that we have towards the circumstances surrounding the uses of words. Much of what we do and say – and even some of our most natural forms of expressions – would be opaque to him. He would fail to understand us when we say, for example, that someone is hot-blooded or that red is a hot colour. In the same vein, even if he were to understand a remark such as “Look, the zookeeper is feeding the lion!”, he would certainly protest when reading Homer’s metaphor “Achilles is a lion”. Compared to ours, his life with words would be heavily constrained by the limits of his ability to appreciate variations in circumstances – and, we might add, by the poverty of his imagination.

However, a limitation on something is not an utter lack of that very thing. Indeed, although confined to the literal, such a speaker would not necessarily be lacking altogether the ability to distinguish contexts in which a term is appropriately used from contexts in which it is not. The metaphorical use of ‘lion’ might very well remain forever unintelligible to him, but we can still imagine him being able – and willing – to distinguish cases where a real lion is in front of him from cases where it is merely a stuffed animal, a cleverly designed prop, or a man in a lion costume. Now, from a distance and seen in a glimpse, the stuffed lion, the prop and the man in costume might very well be indistinguishable from a real lion (after all, the point of an artifice is to look like it is real). However, it is clear that understanding the literal meaning of ‘lion’, implies understanding that a lion is bound to move, to hunt, etc. (i.e. to do things that artifices cannot do). A speaker might not have access to all the relevant circumstances and so might mistakenly call a man in a costume a lion. But if he had access to them (i.e. if he were to see the man getting into the costume,

walking on his two feet before entering into character, etc.) and he meant by ‘lion’ what we usually do by it, then he would recognise that it is inappropriate to call the man a lion – and so that he ought not to do so. The point is that, even with such a speaker for whom what it is correct to say apparently boils down to what is true, the notions of appropriateness and of circumstances remain sufficient to account for his dealing with words. This suggests, I think, that contrary to what Hattiangadi and Whiting assume we are not forced to accept truth and reference among our fundamental semantic notions. These considerations also show that even if we were to reduce our language to its literal use – and to ignore, in the process, some of its essential dimensions –, the possibility of incurring obligations to say something incorrect (on which Hattiangadi’s argument revolves) would not reemerge. Indeed, in failing to see that it might be appropriate in certain circumstances to lie, to use sarcasm, etc., the speaker confined to the literal would go against obligations *we* recognise and abide by, but not against obligations *he* would have himself incurred. That is, despite what it might superficially look like, he does not speak *our* language – in the sense that he does not share what Wittgenstein calls ours form of life. He might use the same words we use, but he does not mean what we do by them. As such, what it is permissible for us to say is not always what it is permissible for him to say. The fact that we do not share the same language – or form of life –, makes it possible for our obligations to diverge.

Now, Hattiangadi – and Whiting – would probably object that if the meaning of an expression is ultimately its conditions of correct use, then these conditions need to be purely semantic (see Hattiangadi, 2006 and 2009 and Whiting 2007 and 2009). Yet, understood in terms of appropriateness, these conditions appear to imply pragmatic considerations, tacit social norms – and even, to some extent, moral imperatives and

epistemic constraints. Of course, I am not denying that; the examples I have discussed are eloquent enough. However, it seems to me that this demand that the semantic should be, in principle, distinguishable from other aspects of language is too demanding. Yes, if meaning is understood in terms of appropriateness, then the lines between the semantic, pragmatic and social aspects of language become more and more difficult to draw. But it is a matter of fact that these aspects are not typically isolable. The use of meaningful expressions tends to imply an attempt to communicate and tends, as such, to involve other speakers and a particular social context. To say that the correct use of an expression implies a sensitivity to the circumstances is precisely a way to grasp this linguistic entanglement.³³

Furthermore, against (1), I think a case can be made that appropriateness is an essentially – and irreducibly – normative notion. Despite what Hattiangadi – and the other leading figures of the Stockholm School – might suggest, distinguishing between the circumstances in which an expression is appropriately used and the circumstances in which it is not is more than simply sorting things into different categories. It is to take a normative stance, to adopt a certain attitude towards the expression. As such, it is not analogous to a mere act of classification. The point is that we, as competent speakers, are not indifferent to how expressions are used. That is why, for example, we tend to correct ourselves when we think we have spoken improperly or manifest our disapproval when we think others are at fault. It also explains the fact that we willingly project expressions into certain new contexts, but not into others. The circumstances surrounding the use of an expression are not merely something of which we are conscious; they are something on which we *act*.

³³ To some extent, the notion of ‘form of life’ plays a similar role when it is invoked to explain what it is to speak a language: it stresses the pervasiveness of language in our life and reveal the constant interplay between what we do and say, what we care about, etc.

Seeing the use of an expression as appropriate – or inappropriate – is ultimately to see it as permissible – or proscribed. As such it is clearly action guiding (i.e. it tells us what we may or ought not to do and say), which, by Hattiangadi’s own admission, makes it normative.

I suspect that Hattiangadi – as well as Glüer and Wikforss – might remain unconvinced. Part of the problem, I think, is that they mistakenly assume that conditions of correct use are to be established independently of what we actually do and say, prior to our actual dealing with words. They see these conditions as standards intelligible, in principle, outside of human life – and so, as essentially disconnected from action.³⁴ And it might not be unreasonable to believe something along these lines if one also assumes, as they do, that conditions of correct use are ultimately conditions of truth. After all, it appears rather uncontentious that whether or not an expression is true of something is independent of what we believe, of our ability to prove it, etc.³⁵ However, Hattiangadi’s position becomes much less appealing when conditions of correct use are understood, as I suggest, in terms of appropriateness. Indeed, the circumstances in which an expression will reveal itself to be appropriate cannot all be given in advance. Every metaphor and projection cannot be anticipated. In dealing with words, new circumstances will always emerge and when that happens, it is clear that we cannot merely refer ourselves to pre-established standards. We are ultimately forced to make a decision as to what to do or say. That ‘hot’ is true only of hot things does not tell us, in itself, that it is appropriate to say of red – but

³⁴ It becomes especially clear when they discuss Whiting’s take on the fairground example. As we have seen both Hattiangadi and Glüer and Wikforss are rather quick to dismiss Whiting’s suggestion that conditions of correct use – like the height requirement – are not only standards, but standards that are *in force*.

³⁵ ‘Hot’ would be true of a hot object even if the object’s temperature was never perceived nor recorded, which suggest that it is possible to stipulate in advance what kind of object ‘hot’ – or any other expression for that matter – will be true of.

not of green – that it is a hot colour or of an impulsive individual that he is hot-blooded. It is the kind of uses we can only sanction as they present themselves. And by sanctioning them – and later by enforcing them –, it is clear that we are doing something normative. In thinking of conditions of correct use as conditions of truth – and so as fixed standards –, Hattiangadi is thus missing something essential about language. Not only is she failing to appreciate the extent to which what we mean by our words emerges from our situatedness, but she also fails to take notice of a genuine kind of normativity that manifests itself in what we do and say.

A New Threat to Naturalism?

My position, in a nutshell, is that conditions of correct use ultimately express our adoption of a normative stance and that they can be seen, as such, as action guiding. Understood in terms of appropriateness or of fittingness with the circumstances, these conditions circumscribe what it is permissible – or proscribed – to do and say. And so, against Hattiangadi, I claim that it is still possible to hold that meaning is normative. The point is that, if a speaker is utterly indifferent to how words are used, if he does not see certain uses as permissible and others as proscribed, then he is not really a speaker. He might utter sounds that look like words – like a parrot can be trained to do – but he does not possess a language.

I would now like to say a few words about the implications of this position for naturalism. At first sight, it is tempting to conclude that my position poses a threat to naturalism. After all, I am suggesting that there is something irreducibly normative about

meaning and since a naturalistic account of meaning cannot make place for it, my position must surely be seen as a charge against naturalism. To this now familiar line of reasoning, I would like to suggest that the matter is not as straightforward. I agree that my position challenges a certain way of conceiving the natural. However, it does so not because it is irreconcilable with any scientific investigation of meaning, but because it forces us to find a place for normativity in nature.

Hattiangadi – like most of the other authors who have contributed to the debate – sees the ‘space of reasons’ as the proper place for normativity and reduces, in the same vein, the ‘space of nature’ to bits of matter and to chains of causal determination.³⁶ That is to say, Hattiangadi already assumes that the ‘space of nature’ is inhospitable to normativity (see Hattiangadi, 2006, 226). For her, if meaning is normative it can thus only be so because, by navigating in the ‘space of reasons’ we come to recognise that we are bound by certain obligations. It makes normativity a matter of determining which standards are justified in guiding what we do and say and which are not. In a word, it reduces normativity to a by-product of rational deliberation.

Hattiangadi says one thing about the fairground example that is, I think, especially enlightening. In response to Whiting’s claim that the height requirement is not only a standard, but a standard *in force*, she insists that “there is a further question whether the standard *ought* to be in force – whether it ought to be accepted and enforced by sanctions.” (Hattiangadi, 2009, 57). To illustrate her point, she proposes a variation on the fairground example. Suppose that there is no height requirement but rather the following standard:

³⁶ Hattiangadi takes both expressions from McDowell and Sellars (see McDowell, 1996 and Sellars, 1997).

R₄: only children who have eaten cornflakes for breakfast on at least one Tuesday in the preceding year are allowed to enter the attractions.

Of course, R₄ could be enforced (employees of the fairground could question children as to what they have eaten before letting them enter the attractions, etc.), but Hattiangadi insists that it would still be an unreasonable standard, a standard that we cannot possibly justify. And for Hattiangadi, this ultimately implies that we cannot have any obligation to abide by it.³⁷

Thus, Hattiangadi's position is that, even if it was enforced, R₄ would not be action guiding since there is no place for it in our 'space of reasons'. I agree with Hattiangadi that this bizarre standard is the source of a certain uneasiness (from whichever angle we look at it, we simply cannot see why a fairground would be compelled to enforce it). However, I think there is a more compelling way to account for this uneasiness than the one she proposes. What we fail to see is the purpose of such a standard. It appears to us arbitrary that what children have eaten – or failed to eat – for breakfast in the past year would have any relevance as to whether or not they can enter attractions. But are we really led to this conclusion because we cannot find reasons for it? I think not. (After all, we will likely reach this conclusion before having the time to reflect on R₄). If we are uneasy with R₄, it is rather because it is too far remote from our form of life, it is because we cannot imagine that someone like *us* would abide by it. It is not to say, of course, that we cannot imagine people for whom what children have eaten for breakfast is of the greatest importance – to

³⁷ Whiting disagrees with Hattiangadi's conclusion, but it is worth noting that he never really questions the underlying picture of normativity on which it rests. Indeed, he is more than willing to concede that meaning or conditions of correct use would be constitutive of the 'space of reasons' if they were normative. As he puts it, they could then be invoked as reasons for employing an expression in one way and not in another. (see for example Whiting, 2009, 537).

the point where it regulates their life. The point is simply that the life of such people would be so distant from ours that we could not recognise ourselves in it. What would be natural for them would be odd for us, etc.³⁸

What I am trying to get at is that, contrary to what Hattiangadi suggests, standards – like the normative stance we adopt towards our expressions – are not just a matter of convention with which we toy in the ‘space of reasons’. They are the expression of our form of life – of our interests, of what is relevant or natural for us, etc. – and as such constantly shape what we do and say. They manifest themselves in our particular way of being in the ‘space of nature’. And if they are normative – as I have argued –, then we need to find a place for normativity in that space. We need to rethink the natural.

Conclusion

In this essay, I have argued that Hattiangadi’s recent objections to the normativity of meaning can be answered if we understand conditions of correct use in terms of appropriateness rather than in terms of truth. Of course, Hattiangadi might be right that truth is not a normative notion – and that we sometimes incur an obligation to lie. I have conceded that much and have agreed with her that both Blackburn and Boghossian’s initial defense of the normativity of meaning and Whiting’s recent attempt to salvage it were unsatisfying because of that. However, I have also tried to show not only that the notions of appropriateness and of circumstances were better candidates than the notion of truth in accounting for conditions of correct use, but also that they revealed something normative

³⁸ These considerations have been inspired in great part by Cavell’s own work on the natural and the conventional (see especially Cavell, 1999, 111-125)

about meaning. Finally, I have suggested that the kind of normativity I have in mind forces us to rethink the natural but is not *de facto* a threat to naturalism.

Conclusion

Rule-Following, Conditions of Correct Use and Obligations

As I have suggested in “Meaning, Understanding and Intelligent Behaviour”, the idea that meaning is essentially and irreducibly normative has been interpreted in at least three ways since the publication of Kripke’s *Wittgenstein on Rules and Private Language*. Some, allegedly sympathetic to insights of the later Wittgenstein, have associated normativism with rule-following (see for example McDowell 1984) while others have rather insisted on the notion of conditions of correct use (see Blackburn 1984 and Boghossian 1989). As for Kripke himself, he seems to have implicitly understood normativism as implying prescriptions (see Kripke 1982). Now, in “Meaning, Understanding and Intelligent Behaviour” as in “The Normativity of Meaning: Rethinking the Natural”, I have tended to treat these three interpretations as fundamentally distinct ways of construing the normativity of meaning – which might have given the impression that I consider them to be somewhat irreconcilable. My position, however, is more nuanced. If I have implicitly accepted a strict demarcation, it is simply because Boghossian and the Stockholm School make a similar assumption. Indeed, they treat rule-following, conditions of correct use and obligations not just as intuitively normative notions, but as distinct – and competing – attempts to make sense of the normativity of meaning. That is why, for example, the arguments they invoke against normativism always focus on one – and only one – of these interpretations. Taken together, they are meant to show that none

of these interpretations is satisfying (i.e. that (i) rules are not constitutive of meaning, that (ii) conditions of correct use are not in themselves normative, and that (iii) meaning does not directly imply obligations). As such, they remain largely silent on the various ways in which rule-following, conditions of correct use and obligations might be interwoven. The impression that each of these interpretations is a distinct account of normativism is also strengthened by Boghossian's and the Stockholm School's understanding of them. As we have seen, they construe rule-following as emerging from our intentions to conform to rules, conditions of correct use as conditions of truth and obligations as action guiding statements. Upon that understanding, there is thus no clear common ground, no explicit conceptual overlap between these interpretations of normativism, which, purportedly, makes it difficult to see how they can relate to one another.

Still, I think the appearances are misleading and that rule-following, conditions of correct use and obligations are best seen as interdependent notions – even if we understand them in the way Boghossian and the Stockholm School urge us to. It becomes rather evident, I think, when we take a closer look at the way their arguments unfold. Take for example their repeated claim that conditions of correct use cannot support normativism. It is made on the ground that these conditions are not in themselves normative and that they do not imply any *obligations*. Or, consider Hattiangadi's attempts to formalise the idea that meaning implies obligations: all of them mention *conditions of correct use*. Finally, Hattiangadi's insistence that obligations are normative because they are action guiding and Glüer and Wikforss's understanding of rule-following in terms of practical reasoning suggest, I think, deep affinities between the two notions.

Moreover, and as I have suggested many times now, nothing forces us to accept Boghossian's and the Stockholm School's understanding of rule-following, of conditions of correct use, etc. We can refuse to recognise, for example, that there is an intentional requirement on rule-following, but still treat rule-following as a useful notion to shed light on the normativity of meaning. This is not to say, of course, that rule-following is an empty notion, that it can be understood in any way imaginable. It is merely to point out that rule-following – like all the other notions that have been invoked to make sense of normativism – refers to something elusive, to something difficult to circumscribe. To follow a rule is not merely to conform to a rule. Someone might momentarily look like he is following a rule without really following it. An automaton might be designed to respond in the same way we do to an order, but it would not make him like us, etc. All these intuitions are constitutive of the notion of rule-following. That much is clear. But sharing these intuitions does not prevent disagreements from arising. Ginsborg's response to Boghossian and Glüer and Wikforss is, on this point, rather telling. Her claim that there is an understanding requirement rather than an intentional requirement on rule-following is not based on her rejection of any of these basic intuitions, but on a different way to approach them. All this suggests that we can be faithful to the notions of rule-following, conditions of correct use and obligations and yet construe them differently. And I suspect that, if we do so – that is, if we distance ourselves from Boghossian and the Stockholm School and come back to the basic intuitions that have given rise to these notions –, their interdependence will reveal itself even more clearly. Ultimately, Boghossian and the Stockholm School's agenda – of offering a systematic criticism of normativism – has a distorting effect: it minimises what

these notions have in common to the point where it becomes difficult to see them otherwise than as competing interpretations of normativism. But we are not prisoners of that agenda.

Salvaging Normativism

It is important to stress that, despite what Boghossian and the Stockholm School sometimes imply, their arguments do not definitively disprove normativism. They do not demonstrate that the very idea of meaning being normative is contradictory, but merely eliminate three ways of accounting for this idea. It is, in part, what I have tried to show in “Meaning, Understanding and Intelligent Behaviour” and in “The Normativity of Meaning: Rethinking the Natural”. We can concede to Boghossian and the Stockholm School that it is impossible to intentionally conform to rules if rules are constitutive of meaning or that conditions of truth are not normative, but that hardly amounts to giving up on normativism. The point is that their criticism of normativism is historically situated. It is a response to Kripke and his first commentators, to the various ways in which these authors understood what it is for meaning to be normative. As such, it does not rule out in advance new accounts of normativism.³⁹ There is thus, we could say, a gap between Boghossian’s and the Stockholm School’s agenda and the actual scope of their criticism. And that is where we can situate Ginsborg’s and Whiting’s attempt to salvage normativism.

It is especially clear, I think, in the case of Ginsborg: unlike Whiting, she never explicitly responds to Boghossian’s or to the Stockholm School’s arguments. Not because she is not aware of them, but because they are for her irrelevant. Indeed, that the intentional

³⁹ That much is rather evident especially if we keep in mind that the very notions that have been historically invoked to account for the normativity of meaning can themselves be reinterpreted.

requirement prevents rules from being constitutive of meaning is without consequence for her own account of normativism – which relies on the very assumption that the intentional requirement can be abandoned in favour of another kind of requirement. Ginsborg’s strategy is not to address the particular challenges to normativism that emerge from Boghossian’s and Glüer and Wikforss’s understanding of rule-following. It is rather to offer a new and more compelling perspective on the basic intuitions that inform the notion of rule-following. In “Meaning, Understanding and Intelligent Behaviour”, I have contended that she is not entirely successful in that endeavour. However, it is for reasons that are foreign to Boghossian and Glüer and Wikforss’s line of reasoning.

As for Whiting, we must recognise that he is more clearly engaged in a debate with Hattiangadi. He explicitly tries to show, on many occasions, that her arguments are flawed or that she is misled by her own analogies. He takes her challenges to normativism seriously and feels compelled to answer them. The reason is that, for him, only – relatively – minor amendments are necessary to salvage normativism. No radical reinterpretation of the terms of the debate is required. It is especially clear, I think, when he maintains, despite Hattiangadi’s contentions, that conditions of correct use are normative. His position is not especially daring: he merely insists that conditions of correct use are standards and that standards have normative implications when they are *in force*. He invokes a basic intuition we have about conditions of correct use but then does little to situate it within a broader understanding of the notion. Still, even in that case, he goes beyond the canonical interpretations of normativism. He does more than echo Kripke and his first commentators. The same goes for when he suggests that meaning is not so much a matter of prescription

as it is a matter of permission. He is ultimately amending Kripke's initial formulation of normativism, even though he remains faithful to many aspects of it.

Thus, both Ginsborg and Whiting attempt to salvage normativism by reinterpreting – or at least clarifying – the notions of rule-following, conditions of correct use and obligations. However, and as I have attempted to show in “Meaning, Understanding and Intelligent Behaviour” and in “The Normativity of Meaning: Rethinking the Natural”, they do not go far enough and ultimately remain in the grip of a problematic picture of normativity. Like Boghossian and the Stockholm School, they implicitly accept that normativity is confined to the inner, that it is something only accessible through the first-person perspective. Indeed, as we have seen, Ginsborg considers the satisfaction of the understanding requirement – and so the very notion of rule-following – to be dependent upon a private sense of appropriateness, while Whiting, following Boghossian and the Stockholm School, sees obligations as emerging from rational deliberation.

In contrast, the account of normativism I have sketched is an attempt to escape this very picture of normativity. In speaking of intelligent behaviour and in suggesting that we are inclined to adopt a normative stance, I meant to draw attention to the many ways in which the normativity of meaning manifests itself in the outer. I wanted to stress that we are not indifferent to how words are used or things are done by pointing out that we often correct ourselves or others, recognise something as a mistake and rectify it, make a decision as to what should be done in new or unusual circumstances, etc. Suggesting that we see purpose in what we – and others – do and say was meant to play a similar role. Indeed, it reveals how the normativity of meaning depends on our sensitivity to the circumstances

and on our ability to cope with them on the spot (like when we amend our response to an order in light of changing circumstances or when we project a word into new contexts).

I must concede that these considerations do not amount yet to a proper account of the normativity of meaning. At best, they are a first sketch. They hint at what distinguishes meaningful expression from mere noise, at what makes our dealing with words essentially and irreducibly normative, but they remain, to some extent, tentative. Still, even at this stage, I think they offer a compelling case for the claim that the normativity of meaning transpires in what we do and say, in our dealing with words – and so that it does not have to be situated behind the façade of behaviour or outside of nature, in the space of reasons. I remain unclear as to what precisely distinguishes intelligent and mechanical behaviour, being even inclined to believe that it might not always be possible to draw a clear line between the two. I am also unsure as to what should ultimately count as adopting a normative stance or seeing purpose in something. I have offered some examples, but clearly more needs to be said. Nonetheless, I remain convinced that these considerations offer a promising avenue to found the normativity of meaning on new grounds – and to rethink our basic intuitions about rule-following, conditions of correct use and obligations.

Moving Forward

In “Meaning, Understanding and Intelligent Behaviour” and in “The Normativity of Meaning: Rethinking the Natural”, I have suggested that my position in the debate surrounding normativism is in line with many philosophical insights of the later Wittgenstein and of Cavell. I have referred, for example, to Wittgenstein’s discussion of

the language-game of the builder and the assistant or to Cavell's appropriation of the notion of form of life, but all this remains rather superficial. The affinities between the account of normativism I have sketched and Wittgenstein's and Cavell's writings are, I believe, more far-reaching.

The claim that to use an expression meaningfully is to see the purpose of the expression and to adopt a normative stance toward it has many implications. One of which is that its meaning cannot be circumscribed once and for all by a set of rigid rules. Indeed, as I have suggested in "Meaning, Understanding and Intelligent Behaviour", we are not at a loss when we face new – or unanticipated – circumstances (that is why, ultimately, we are not behaving merely mechanically but intelligently). Of course, at one point, we might have abided by certain rules – and have taken them to be quite definitive. However, if they were to reveal themselves to be useless, by remaining silent in particular circumstances as to what is to be said or done, we would be able to go beyond them. We would be able to decide if the expression has taken on a new purpose, if it can be appropriately projected, etc. And ultimately, that is why our words are not mere noises: our life with them cannot be reduced to blindly abiding by fixed rules.

Now, I think Wittgenstein suggests something similar in an often-neglected remark of the *Philosophical Investigations*. Comparing language to games, he goes on to say:

We can easily imagine people amusing themselves in a field by playing with a ball like this: starting various existing games, but playing several without finishing them, and in between throwing the ball aimlessly into the air, chasing one another with the ball, throwing it at one another for a joke, and so on. And now someone says: The whole time they are playing a ball-game and therefore are following definite rules at every throw.

And is there not also the case where we play, and make up the rules as we go along? And even where we alter them – as we go along (Wittgenstein, *PI*, 83)

This remark just precedes the so-called rule-following considerations – on which Kripke and many other commentators of Wittgenstein have focused. It should thus be read, I think, as a word of caution about that notion of rule-following. Indeed, as I see it, it suggests something along these lines: in our dealing with words, following a rule is not just a matter of abiding by a rule. It also requires us, in certain circumstances, to amend the rule, to go beyond what the rules say. As such, to use an expression meaningfully is not akin to making a move in a game of chess (where the rules settle once and for all which moves are permissible and which are not), but rather to playing with a ball in a field (where rules can be abandoned or created, variously altered, etc.)

Additionally, the claim that the normativity of meaning is not something hidden shares many affinities with Wittgenstein's and Cavell's broader position on the inner/outer distinction. To begin with, my contention with Ginsborg's suggestion that understanding amounts to a feeling of appropriateness, to something essentially private closely parallels Wittgenstein's repeated denial that understanding consists in something 'going on' in the mind. For him, whether or not someone understands an order, a chart or a signpost does not depend on the presence or absence of a feeling of understanding, but on what he does in response to the order, etc.⁴⁰ That is, for Wittgenstein, one's understanding reveals itself in one's behaviour and is thus, in principle, accessible to the third person perspective. Moreover, and more generally, Wittgenstein insists, on many occasions, that 'mental' notions (such as meaning, intention and even pain) carve up not the inner, but rather parts

⁴⁰ See for example Wittgenstein's discussion of what it is to understand the rule '+2' (see especially Wittgenstein, *RP*, 152 and 154).

of the outer. We see it, for example, when he suggests that the difference between ‘feeling’ pain and feigning pain is not to be found in any feeling that would accompany pain-behaviour in the former case –but not in the latter (see for example Wittgenstein, *PI*, 304). Indeed, the mistake, for him, is to think that ‘feeling’ pain and feigning pain are undistinguishable from the third-person perspective, that they unfold in the exact same way in the outer. Of course, if they were identical patterns of behaviour, then only something inner could differentiate them, but they are not (see Wittgenstein, *LW2*, 42).

The case of pain is, I think, rather representative of Wittgenstein’s position toward the mental and the inner/outer distinction. Moreover, it shows quite clearly how Wittgenstein’s insights on the matter relate to my present philosophical concerns. Indeed, there is an undeniable parallel between the worry that we might never really know if someone is in pain and the worry that we can never have the assurance that someone is following a rule – rather than merely conforming to it by accident. Boghossian, Glüer and Wikforss and Ginsborg, in implicitly accepting that the distinction between following a rule and merely conforming to it depends on something hidden behind the façade of behaviour, remain unable to satisfyingly deal with such a worry. That is why, for example, they are bound to take seriously the possibility that we might just be like the parrot or the automaton. They fail to appreciate the many nuances in our behaviour that distinguish us from them – very much like the philosophers who see pain as something inherently private fail to appreciate the many ways in which being in pain is unlike feigning it.

Finally, both Wittgenstein and Cavell often insist that what we mean by our words and what we do with them is rooted in our particular form of life (see Cavell, 1999 and 2002). As such, it reflects our cultural upbringing, of course, but also our interests,

inclinations, physiological features, etc. That is what Wittgenstein suggests, I think when he claims in the *Philosophical Investigations* that “giving orders, asking questions, telling stories, having a chat, are as much a part of our natural history as walking, eating, drinking, playing” (Wittgenstein, *PI*, 25). Now, when I claimed in the second article that we need to rethink the natural, I had in mind something along these lines: adopting a normative stance and seeing purpose is constitutive of our form of life. It constantly informs our way of being in the world and as such, reveals itself in what we do and say.⁴¹ Because we adopt a normative stance and see purpose, we do not inhabit the space of nature like a parrot or an automaton.

My goal in the two preceding articles was to situate myself in the recent debate surrounding normativism. I often mentioned Wittgenstein and Cavell, but I never hide behind their authority. I treated their writings as a starting point, from which I could depart, if necessary. Still, in conclusion, I wanted to suggest that it might be useful to come back to their writings yet again to further clarify the notions of intelligent behaviour and of normative stance I have introduced. I do not think my account of normativism is an unoriginal repetition of what Wittgenstein and Cavell have already said, but it is certainly a way to inherit their philosophical insights.

⁴¹ That is why I have suggested that a speaker confined to the literal would neither speak the same language as us, nor share our form of life.

Bibliography

- Blackburn, S. (1984). The Individual Strikes Back. *Synthese*, 58, 281-301.
- 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.
- Bridges, J. (2014). Rule-following Skepticism, Properly So Called. In *Varieties of Skepticism: Essays after Kant, Wittgenstein, and Cavell*. A. Kern and J. Conant (eds.). Berlin: de Gruyter.
- Bykvist, K. and Hattiangadi, A. (2007). Does thought imply ought?. *Analysis*, 67 (4), 277-285.
- Cavell, S. (1999). *The Claim of Reason* (New Edition). Oxford: Oxford University Press.
- Cavell, S. (2002). *Must We Mean What We Say?* (Second Updated Edition). Cambridge: Cambridge University Press.
- Cavell, S. (2013). *This New Yet Unapproachable America*. Chicago: The University of Chicago Press.
- Ginsborg, H. (2010). Review of *Oughts and Thoughts: Rule-Following and the Normativity of Content*, by Anandi Hattiangadi. *Mind*, 119, 1175-1186.
- Ginsborg, H. (2011). Primitive Normativity and Skepticism about Rules. *The Journal of Philosophy*, 108 (5), 227-254.
- Ginsborg, H. (2012). Meaning, Understanding and Normativity. *Proceedings of the Aristotelian Society Supplementary Volume*, 86, 127-146.
- 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.

- Goldfarb, W. (1983) I want you to bring me a slab: Remarks on the opening sections of the *Philosophical Investigations*, *Synthese*, 56, 265-282.
- 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.
- Kripke, S. (1982). *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Cambridge, MA: Harvard University Press.
- McDowell, J. (1984). Wittgenstein on Following a Rule. *Synthese*, 58, 325-363.
- McDowell, J. (1996). *Mind and World*. Cambridge, MA: Harvard University Press.
- McGinn, C. (1984). *Wittgenstein on Meaning*. Oxford: Blackwell.
- Mulhall, S. (2007). *Philosophical Myths of the Fall*. Princeton: Princeton University Press.
- Pettit, P. (2002). *Rules, Reasons and Norms*. London: Oxford University Press.
- Rhees, R. (1959). Wittgenstein's Builders, *Proceedings of the Aristotelian Society New Series*, 60, 171-186.
- Sellars, W. (1997). *Empiricism and the Philosophy of Mind*. Cambridge, MA: Harvard University Press.
- Whiting, D. (2007). The Normativity of Meaning Defended. *Analysis*, 67 (2), 133-140.
- Whiting, D. (2009). Is Meaning Fraught with Ought?. *Pacific Philosophical Quarterly*, 90, 535-555.
- Whiting, D. (2016). What is the Normativity of Meaning?. *Inquiry*, 59 (3), 219-238.
- Wikforss, A. (2001). Semantic Normativity. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 102 (2), 203-226.
- Wittgenstein, L. (1967). *Zettel*. G.E.M. Anscombe and G.H. von Wright (eds.). Berkeley, Los Angeles, and London: University of California Press.
- Wittgenstein, L. (2009). *Philosophical Investigations* (Fourth Edition). Oxford: Wiley-Blackwell.
- Wittgenstein, L. (1993). *Last Writings on the Philosophy of Psychology: The Inner and the Outer, Volume 2*. G.H. von Wright and H. Nyman (eds). Oxford: Blackwell.
- Wright, C. (1989). Critical Notice of Colin McGinn's *Wittgenstein on Meaning*. *Mind*, 98, 289-305.