

# Inferentialism and the Normativity of Meaning

Jaroslav Peregrin

Received: 4 February 2010 / Revised: 27 August 2010 / Accepted: 29 August 2010 /  
Published online: 14 September 2010  
© Springer Science+Business Media B.V. 2010

**Abstract** There may be various reasons for claiming that meaning is normative, and additionally, very different senses attached to the claim. However, all such claims have faced fierce resistance from those philosophers who insist that meaning is not normative in any nontrivial sense of the word. In this paper I sketch one particular approach to meaning claiming its normativity and defend it against the anti-normativist critique: namely the approach of Brandomian inferentialism. However, my defense is not restricted to inferentialism in any narrow sense for it encompasses a much broader spectrum of approaches to meaning, connected with the Wittgensteinian and especially Sellarsian view of language as an essentially rule-governed enterprise; and indeed I refrain from claiming that the version of inferentialism I present here is in every detail the version developed by Brandom.

**Keywords** Meaning · Normativity · Inferentialism

## Introduction

There may be various reasons for claiming that *meaning is normative*, and additionally, very different senses attached to the claim. However, all such claims have faced fierce resistance from those philosophers who insist that meaning is not normative in any nontrivial sense of the word—*viz.* Glüer and Pagin (1999), Wikfors (2001), Boghossian (2005), Hattiangadi (2006), Glüer and Wikfors (2009) and Hattiangadi (2009). In this paper I will sketch one particular approach to meaning

---

J. Peregrin (✉)  
Institute of Philosophy, Academy of Sciences of the Czech Republic, Prague, Czech Republic  
e-mail: jarda@peregrin.cz  
URL: jarda.peregrin.cz

J. Peregrin  
Faculty of Philosophy, University of Hradec Králové, Hradec Králové, Czech Republic

claiming its normativity and defend it against the anti-normativist critique: namely the approach of Brandomian inferentialism.<sup>1</sup> However, my defense is not restricted to inferentialism in any narrow sense for it encompasses a much broader spectrum of approaches to meaning, connected with the Wittgensteinian and especially Sellarsian view of language as an essentially rule-governed enterprise; and indeed I refrain from claiming that the version of inferentialism I present here is in every detail the version developed by Brandom.

What I take as characteristic of this approach is the view that meanings are roles which are acquired by types of sounds and inscriptions in virtue of their being treated in accordance with the rules of our language games, roughly in the same way as wooden pieces acquire certain roles by being treated in accordance with the rules of chess. It follows—*inter alia*—that (i) a meaning is not an object labeled (stood for, represented ...) by an expression; and that (ii) meaning is normative in the sense that to say that an expression means thus and so is to say that it ought to be used in a particular way. The inferentialism of Brandom (1994) represents the most systematic elaboration of this train of thought.

### Inferences vs. Inferential Rules

Let me start with a brief survey of those features of the Sellars's and Brandom's approach that will be relevant for our case. The first thing to realize is that the key concept of the kind of inferentialism I am going to discuss here is that of *inferential rule*, not *inference*. Neither Sellars, nor Brandom are interested in what happens within an individual mind; what they are after are rules as socio-linguistic matters.<sup>2</sup>

Inferential rules, of course, concern what we should do when we communicate; but in fact it may be helpful *not* to see them as telling us what to do, for this may be potentially severely misleading. They dictate what is *allowed*; or (the other side of the same coin) what is *prohibited*. Therefore we should not see them as *commands*, but rather as *constraints*.<sup>3</sup> An inferential rule does not tell us directly what to do (save in the sense of dictating what we are *allowed* to do), but rather what *not* to do. To say that *A* is inferable from  $A_1, \dots, A_n$  is *not* to say that whoever asserts (thinks, ...)  $A_1, \dots, A_n$ , should also assert (think, ...) *A*; but rather that it is incorrect to assert  $A_1, \dots, A_n$  and to *deny* (i.e. to preclude the possibility of asserting) *A*.

<sup>1</sup> I have dealt with other kinds of objections to inferentialism elsewhere (see Peregrin 2009).

<sup>2</sup> This is important to keep in mind, for besides this *normative* variety of inferentialism we can also consider a *causal* variety (which, in effect, is a subspecies of functionalism, well-known from the philosophy of mind); and the latter is very often mistaken for the former. (Unlike the normative variant, the causal variant concentrates on the role of an expression within inferences actually carried out by speakers or thinkers.) It is the latter which some philosophers appear to have in mind when they talk about *inferential role semantics* (Boghossian 1993; Fodor and Lepore 1993); but I will reserve the word *inferentialism* for the normative version (and it was coined, by Brandom, in this very sense).

<sup>3</sup> Von Wright (1963) uses the terms *command* and *prohibition* in a similar sense; while Sellars (1974, 422) uses *constraint* and *incentive* and claims: "It should be stressed that the uniformities involved in meaningful verbal behavior include *negative* uniformities, i.e. the avoidance of certain combinations, as well as *positive* uniformities, i.e. uniformities of concomitance. Indeed, negative uniformities play by far the more important role, and the rules which govern them are to be construed as *constraints* rather than incentives."

What is the status of the rules? How do they exist? One answer might be that they are *regularities* of behavior, and that *following a rule* amounts to *being in step with other language users*. It is important to realize that this is *not* the sense entertained by the kind of inferentialism I am defending. Using a language presupposes, according to inferentialism, a kind of a *normative attitude* to the utterances of others (as well as to my own). I *take* what they do for *right* or *wrong*, which manifests itself by (though is not reducible to) my ‘rewarding’ those who do the right things, and ‘sanctioning’ those who do not (the ‘reward’ or ‘sanction’ would likely be my bestowing upon them a *status*—ranging, say, from ‘reasonable speaker’ down to ‘notorious babbler’—, which may, if resonating with their evaluation by other people, become their ‘official’ status within the community in question).

In the simplest case, a rule exists via a symbolic articulation (statement, command, instruction ...). But this is not how rules exist in general—at least in some cases rules must exist without being explicit, being carried merely by the normative attitudes. The reason for this is that a symbolic articulation can act as a rule only for those able to *interpret* it. Even a signpost, as Wittgenstein (1953, §85) pointed out, must be interpreted. However, interpretation is precisely the kind of enterprise that requires rules. Thus, to follow an explicit rule one needs a further rule, and hence if every rule were to be explicit, an infinite regress would be looming.

As for the normative attitudes, they cannot be seen as something static, furnishing human actions, once and for all, with clear-cut labels “correct” and “incorrect”. These attitudes, upon being made explicit, themselves become subject to our “game of giving and asking for reasons” (Brandom’s 1994, term) and thus are continually subject to possible amendments in the light of reasons. Hence, though each of such attitudes is *authoritative* insofar as the (*in*)*correctness* that it intends to bestow on some actions is construed as something permanent and unconditional, it nevertheless includes, as its integral component, an openness towards the possibility of its *re-evaluation*. And due to the permanence integral to the concept of correctness, any such re-evaluation comes to be understood as *the discovery of an error*—hence any correctness judgment, however authoritative, is itself subject to future evaluation with respect to *its own correctness* and is thus, in this sense, never definite.

The rules of language can be seen as furnishing individual sentences, and consequently individual words, with *roles*, which then, from the inferentialist viewpoint, appear as meanings. Hence, not every inference is constitutive of meaning (in fact none is!)—only inferential *rules* are so constitutive. It is not the case that expressions have their inferential roles *as a result of* having certain meanings; having the inferential role *is* having the meaning.

Let me note that from the viewpoint of Brandom, the picture of inferentialism, as drawn so far, remains ‘superficial’ in that it has only envisaged the ‘surface’ layer of the functioning of rules. In this paper we intend to stay mostly on this very level, so we will not ply the underlying level here. The only thing I wish to stress is that the mechanisms animating the surface level are *essentially social*—concepts such as *rule* and *inference*, *sanction* and *reward*, etc., are essentially a matter of social interaction; and if we do not see them as such, we are inevitably missing the point of this kind of

enterprise. This also leads to the conclusion that meaning is to be sought seeing language as a *social institution*, rather than as, say, a psychological or biological reality.<sup>4</sup> This reaffirms the point made above, that the variety of inferentialism I am talking about is not (directly) about any causal proceedings within the human mind/brain. The proceedings on which the inferentialist concentrates are those of a certain *social reality* (i.e. a kind of reality exemplified by such entities as NATO, university positions etc.—a reality which does not exist apart from human attitudes, but which is objective in the sense of enjoying independence from the attitudes or will of any individual human).

### Should We Assert only the Truth?

A common objection to the normativity of meaning is based on the claim that semantics can yield normative consequences only via a supplementary norm, such as *One ought to assert that p only if (she is convinced that) p*. To give an example of this kind of objection, let me quote one of the most influential exponents, Paul Boghossian (2005, 207):

If there is to be an interesting thesis of the normativity of meaning, we ought to be able to derive a should or an ought from the mere attribution of meaning to someone and without having to rely on any auxiliary desires that that person may or may not have. But can we do that? Does it follow from the mere fact that I mean addition by ‘+’ that I should not lie or mislead? There may be, for all I know, a *moral* prohibition against lying or misleading; but are there such prohibitions flowing from the nature of meaning itself?

Boghossian’s (2005, 212) verdict is then unequivocal:

To put the matter concisely, the linguistic version of the normativity thesis [i.e. the thesis *One ought to believe that p only if p*—J.P.], has no plausibility whatever; and the reason is that it is not a norm on assertion that it should aim at the truth, in the way in which it is a norm on belief that it do so. Thus, the only imperatives that flow from attributions of linguistic meaning are hypothetical imperatives.

Hence what Boghossian claims is that while there is a norm that our *beliefs* “should aim at the truth”, there is no norm that our *assertions* should do so. I think otherwise; in fact I think that while there *cannot be* a norm of the former kind, there *is*—in a sense—one of the latter kind. As for believing the truth: as it seems to me to

<sup>4</sup> According to Brandom (1994), the basic gears of the underlying communal machinery are the deontic statuses of *commitment* and *entitlement*. This way of approaching the speech acts leads to a kind of pragmatics which is essentially normative: it characterizes the speech acts in terms of the kinds of rules that govern them and in terms of those changes of normative statuses of the participants of communication which they bring about. Participation in linguistic communication essentially involves scorekeeping. Semantics, then, is in effect *nothing else* than a theory of roles conferred on linguistic tokens by the rules, i.e. of the ways in which playing these tokens is capable of changing the deontic statuses of the player and her companions.

be obvious that one is not free to decide what to believe,<sup>5</sup> we cannot say that one's belief ought to "aim at truth", simply because *ought* implies *can* (and of course also *need not*).<sup>6</sup> (But this is not something I want to discuss here in detail.) And as for asserting the truth: I follow Sellars (1992, p. 101) in understanding truth as correct assertability, which entails that an assertion is *correct* if what is asserted is true; and hence, in this sense, that our assertions *ought to* "aim at truth" (though, of course, as noted in "[Inferences vs. Inferential Rules](#)", this does not mean that we are *obliged* to assert *everything* that is true, but rather that we are *forbidden* to assert *anything* which is *not* true, i.e. that we are *permitted* to assert *only* the truth).

How does it square with the obvious fact, pointed out by Boghossian, that there are situations when we cannot but say that one ought to lie? Well, there are always various kinds of rules one is bound by and hence various sources of "oughts" which may be in conflict. Of course I can imagine a situation where a moral obligation might dictate me to violate the rules of chess—but does it imply that the rules are non-existent or that they do not tell me how I *ought to* play chess? And should an analogous conflict tell me that the rules of language constitutive of the meaning of an expression do not tell me how I should use the expression?

Although the consequence of the inferentialist standpoint is that one ought not to assert what is not true (*pace* Glüer and Wikforss 2009, 38), it is crucial to realize that this is not a pragmatic principle over and above semantic principles. For an inferentialist, truth is nothing more than a nickname for a status certain sentences have *vis-à-vis* the rules of the game of giving and asking for reasons; hence the result that one ought to assert only true sentences is nothing over and above unpacking the concept of truth.

Consider a parallel argument against there being a norm of asserting the truth by Wikforss (2001, 205–6):

Consider the case where I misperceive and utter 'That's a horse' of a cow. What semantic norm do I then violate? I see the animal, believe it to be a horse and, consequently, utter 'That's a horse'. Although I have made a false judgment, I have not broken any semantic norms.

But suppose I decide to call the correct chess moves *cool* and I say: "Suppose I misperceive and move a rook as if it were a bishop. What rule of chess do I then violate? I see the piece, believe it to be a bishop and, consequently, move it diagonally. Although I have made a move that is not cool, I have not broken any rule of chess." To avoid misunderstanding, I do not pretend to have presented an argument for the claim that truth *is* correct assertability; and there is not enough room to present it here (hence I must refer the reader to Sellars 1992, Chapter IV). I just invite the reader to admit this as a possibility. The argument I offer is to challenge the criticism that there *cannot* be a norm requiring us to assert the truth (though, to repeat, I am certainly not going to claim that we would be required to

<sup>5</sup> Cf. the classic text of Williams (1973).

<sup>6</sup> We *can* say that it is a norm that we *ought to interpret our peers as believing* the truth. This is the celebrated Davidsonian principle of charity; but this is a far cry from claiming that we *ought to believe* the truth.

assert the truth in the sense of asserting everything that is truth, but rather in the sense of asserting only what is true).

In any case, the inferentialist analysis of meaning reveals a *deeper* source of normativity than an alleged norm that we should assert the truth: meaning, according to the inferentialist, is normative in the sense that when I say that an expression means thus and so, then what I say does not amount to stating a fact, but rather invoking a propriety: it is stating that the expression is *correctly* used thus and so. (Hence the inferentialist rejects Boghossian's claim that to derive "a should or an ought from the mere attribution of meaning" we need an additional premise, such as his "normativity thesis".) True, on one of its readings, this may still be read as stating a kind of fact, namely that an activity or a community is—as a matter of fact—governed by certain rules (hence on this reading we treat the propriety in question *as a fact*); however, there is a second and crucial reading on which this is not the case, for the claim does not amount to a declarative statement at all, it is rather an *endorsement*. As Sellars puts it in his letter to Chisholm,<sup>7</sup>

My solution is that

' . . . ' means ---

is the core of a unique mode of discourse which is as distinct from the *description* and *explanation* of empirical fact, as is the language of *prescription* and *justification*. (Chisholm and Sellars 1958, p. 527)

Is it true then that, as Boghossian claims, "the only imperatives that flow from attributions of linguistic meaning are hypothetical imperatives"? Well, if we read the attributions of meaning in the first of our two ways, then it is: it follows merely that you should use the expression thus and so *if you want to speak a certain language, or to belong to a certain community*. But on the second reading, and this is crucial, the qualification drops out, it follows simply that the expression should be used thus and so (full stop).

I suspect that Boghossian would object that this kind of normativity is still not "genuine" and that hence meaning is not normative in the "genuine" sense of the word he has in mind, for the fact that in some situations we feel compelled to say that someone ought to use an expression at variance with its ('standard') meaning (e.g. that one ought to assert a false sentence). This would suggest that a norm can be "genuine" only if it is *indefeasible*. But as far as I can see, the defeasible/indefeasible opposition is an opposition *different* from the hypothetical/categorical one—even the most essential moral or legal norms may sometimes conflict with one another and hence are not immune from being discharged. If the only genuine norms were indefeasible ones, then I would see no genuine norms around at all. We know all too well that for *any* rule, including the most categorical ones, there are situations when we take it for discharged—think about the Biblical *thou shalt not kill* and all the *except*'s that have come to accompany it during the history of our civilization!

<sup>7</sup> To avoid misunderstanding: as I stressed that at least some rules of language must remain merely implicit, let me now stress that of course the articulation of what I call normative may amount merely to making the normativity explicit, not necessarily to instituting it. To institute it necessitates taking up normative attitudes—especially towards those who thwart a rule: correcting them, deterring them from doing what they do, denigrating or punishing them. To say that an expression means thus and so is already to move on a metalevel w.r.t. the original game and is not essential for the presence of the normativity.

I am fond of describing the situation in terms of an ‘internal space’ that some systems of rules have the ability to constitute. From outside of the space we can only report the fact that the rules are in force for the insiders, but once we join them, they start to be in force *for us* and hence be in force (full stop); and claiming this does not amount to stating a fact, it is a different speech act. Let me call the former reading the ‘outsider’ reading and the latter the ‘insider’ one.

Claims to the effect that something is correct or that something ought (not) to be done (I will call them *normatives*, for short), on the insider reading (I will call them *true normatives*), are speech acts distinct from assertions or reports. They do not report that something is (not) the case, they point out that something *ought (not) to be* the case, hence they always involve the utterer’s taking a rule as being in force, her endorsing it. In this respect, they are similar to oaths of loyalty: they always involve one’s decision to assume a certain status, namely to bind oneself by a rule, and in this sense they *institute* something (namely a certain social link) rather than *report* it. However, the case when the institution happens in a single instant (as in the case of signing an oath) is only a very special case; more generally, binding oneself with a rule is more like the case of loyalty that is not formally established with an instant oath, but is continuously testified by one’s performances and declarations. Normatives of this kind involve the instituting and upholding—or, as the case may be, amending or contravening—a rule.

This is, of course, not the only thing that normatives do: besides this, they may express that the rule in question, as applied to a particular case, renders the case right or wrong. If I tell you “You should not kill this cat” I claim that (given a certain rule to which I, and presumably you, submit) killing this cat would be wrong. Concentrating on this, we might say that the normatives report a specific kind of fact, *viz.* normative facts—but the possibility of taking this characterization literally is compromised by the alleged fact being a fact only by virtue of a rule that does not exist totally independently of the statement, for the statement takes part in its constitution.<sup>8</sup>

Hence compare the following two claims:

- (1) Killing this cat would be easy
- (2) Killing this cat would be wrong

Both of them can be seen as classifications: they classify a certain hypothetical action from a certain viewpoint. But whereas (1) uses a classificatory criterion that is wholly independent of the classification and can thus be read as objective in the most straightforward sense, things are different with (2). Though it can perhaps be read in the same objective and hence disengaged way (the ‘outsider’ reading), the important point is that it can also be read in an alternative way, where its aim is not only to classify, but at the same time to uphold the criterion that is employed, to declare one’s allegiance for it (the ‘insider’ reading). Hence two sentences, though their grammatical structure is the same, may be used to accomplish dissimilar speech acts.

<sup>8</sup> Remember that we are talking about the insider reading of the normatives. On the outsider reading, the constitutive ingredient of the claim falls out and the claim may be read as a report much more straightforwardly.

This should clarify why an inferentialist disagrees with those philosophers who claim that normatives are doomed to be dubious unless their content is presented in a non-normative idiom. The content of true normatives is of a kind different from that of declaratives—it is better to think of them on a par with interrogatives or imperatives. The inferentialist is able and willing to explain how human communities came to entertain them, and to explain what they are used for—but not to translate them into a non-normative idiom.

### Do Inferential Rules Guide Us?

We have stated that rules of language constitute meanings of words just like the rules of chess constitute the ‘values’ of pieces. But does this not rob them of being *rules* in the fully-fledged sense of the word, namely of being something that *guides us in our behavior*?

Glüer and Pagin (1999) argue for the thesis that “rules that can determine meaning ... are *not* capable of guiding speakers in the ordinary performance of speech acts”. They base their argument on the distinction between what Searle (1969, 34) calls *constitutive* and *regulative* rules:

Regulative rules regulate a pre-existing activity, an activity whose existence is logically independent of the rules. Constitutive rules constitute (...) an activity the existence of which is logically dependent on the rules.

Glüer and Pagin then argue that constitutive rules cannot enter into practical syllogisms which guide us in our actions, and hence insofar as some rules are meaning-constitutive, they cannot be action-guiding.

What I take to be the most important points of Glüer and Pagin are the following theses:

1. “A practice is constituted by a set of rules if it is possible to engage in that practice *only* insofar as the rules of that set are in force for the agent.” (221)
2. We should not think that a speaker must actually *follow* a meaning-constitutive rule in order to express the meaning constituted by it; “what is decisive is precisely that the rule is *in force* for the speaker.” (222)
3. “Rules that *can* determine meaning, ... i.e., rules that can be regarded as *constitutive of* meaning, are not capable of guiding speakers in the ordinary performances of speech acts.” (207)

I think all these points are essentially correct. But I find it imperative to point out that the last thesis is correct *only in one possible sense of “guiding”*—a very narrow sense. It is correct in the sense that meaning constitutive rules are not what I have already dubbed *commands*: at no (or almost no) point of our language game do they tell us what to do.

Does this mean that norms do not enter such deliberations at all? I think they do, but their role is more delicate. Consider a chess player pondering her next move: she considers checking the opponent’s king and is weighing up her options. “Maybe I could move my rook diagonally? No, that would be an illegal move ...” What happens here? If we look at her as not playing chess, but merely moving pieces of



wood, we can see her as reasoning using the rules of chess as premises. (*I could move this piece diagonally. But this would be an incorrect thing to do. Hence I will not do it.*) However, as a chess player, moving *chess* pieces, this kind of reasoning is not available to her: the rules cannot be used as premises of reasoning, for they are what forms the space within which the *rook* becomes what it is. Hence the level on which the practical deliberation may be seen as taking place is buried under the level which meets the eye—the level on which the pieces already are *rooks*, *pawns* and *kings*.

Using language can be looked at similarly. We can look at a speaker as weighing up which kind of sound she should emit, using the rules of the current language game as the premises: “Maybe I could voice the sound *The animal over there is a dog?* But I have just voiced *The animal over there cannot be a mammal*, which is not compatible with it ...” On the surface level, where we see her as already a competent speaker uttering meaningful sentences, such patterns of reasoning lie buried within the assumption that the words and constructions employed mean what they do.

Hence it would seem that there *are* patterns of practical reasoning in which rules of the meaning-constitutive kind *can* play the role of premises. But they are, usually and characteristically, buried deeper within the process—they are *constitutive* of specific kinds of reasoning. Only because the rules of chess exist in the first place do I even think of wanting to *check a king*; and only insofar as there exist rules of language am I at liberty to consider wanting to *say that the animal over there is a dog*. As I have written elsewhere,<sup>9</sup> the rules open up a brand new space of possibilities.

It is, I am convinced, badly misleading to think about these new possibilities as merely new guises for some old ones. This was part and parcel of the treatment of rules due to Searle (1969), discussed (critically) by Glürer and Pagin. Searle proposed that what the rules achieve in general is to make an old action count as a new one. Thus, moving a certain piece of wood in a certain way comes to count as checking the opponent’s king, and emitting a certain sound comes to say that an animal is a dog. However, I think that to see things in this way is fundamentally misleading. What is it that comes to count as checking the king? What kind of move with what kind of object? We know that the piece need not be wooden, and, moreover, we know that we do not necessarily need *any* tangible piece, we can play chess and check merely by declaring certain letters and numbers. And similarly, we know that to say the animal over there is a dog does not require us to emit any clearly specified kind of sound; we can do it in many direct and indirect ways!

The point, as I see it, is that rules of this constitutive kind may constitute wholly *new* kinds of actions which are not reducible to the old ones, i.e. to the ones which were available before the rules were established. Hence the rules do not merely restrict the spectrum of possibilities we can choose from, they bring these possibilities to life in the first place. Thus, if Glürer and Pagin wish to claim that “the problem with constitutive rules ... is that there simply isn’t anything that they require” (217), this is true in one sense only—there isn’t anything they require in the *command* sense of “require”; however, there is something they require in the *restrictive* sense. They gerrymander our playground, and by doing this they create a new space.

<sup>9</sup> See Peregrin (2010b).

To what extent, then, do the conclusions of Glüer and Pagin contradict inferentialism as exposed here? I agree that the crucial role of normatives within practical reasoning is not as premises of patterns that lead to actions. (Though I think that there may be patterns having normatives as premises, *viz.* the above one, I do not think they play an important *explicit* role.) However, this does not preclude norms (be they made explicit in terms of normatives, or not) from playing a fundamental role in semantics and hence in meaningful speech—on the contrary, they are *constitutive* of the very realm of semantics. Hence, insofar as Glüer and Pagin claim merely that normatives do not guide us in the command sense of guiding, we have no quarrel.

### Are Inferential Rules Prescriptive?

There is a very simple way to articulate the reasons leading the inferentialist to the normative construal of meaning: (i) there are correct and incorrect ways we can use our expressions; hence (ii) there are rules governing the usage of the expressions; and hence (iii) the meaning of the expressions are roles *vis-à-vis* the rules: Just like what makes a piece of wood into a chess rook is that it is treated according to a certain set of rules, what makes a sound- or inscription-type into an expression meaning thus and so is that it is treated according to a certain other set of rules.

The objection discussed in the previous section concerned the step from (i) to (ii): the claim was that if there are rules of language, they are not rules capable of “guiding us” in our linguistic practices. Our conclusion was that this is true on an all too narrow construal of *guiding*, but false on a wider construal, which still enables us to talk about normativity here. More radical challenges to this step deny this possibility altogether, and hence deny that the step from (i) to (ii) can be legitimate: there are no rules, so the objection goes, governing our usage of expressions, or at least not rules that would be *prescriptive* in a sense capable of justifying a normative construal of meaning. Thus Glüer and Wikforss (2009, 36) write:

If I mean green by ‘green’, then ‘green’ is true only of green things, and if I say ‘That is green’ while pointing at a red object, I have said something false. But it does not immediately follow that I have failed to do what I ought to do—even from a merely semantic point of view. There are non-normative uses of ‘correct’, and this is one of them. The relevant notion of correctness in this context is that of semantic correctness. ... Semantic categorization is non-normative in precisely this sense: it has no direct normative consequences.

I find this argument puzzling. As far as I can see, the words *correct* and *ought* to are closely interrelated simply by way of their semantics—I think that how we normally understand the classification of an action as *correct* is precisely that it *ought to be carried out* (and classifying a state as correct as something that *ought to be brought about*<sup>10</sup>). Maybe there are reasons to reject this intuition and accept, together with the authors, that “there are non-normative uses of ‘correct’”—but they

<sup>10</sup> For a distinction between what ought to be done (an *ought-to-do*, in Sellars's term) and what should be brought about (an *ought-to-be*) see Peregrin (2010a).

do not give any such reasons, and rather shift the burden of proof to their opponents (*ibid.*, footnote 8).

[T]he normativist might feel safe in just digging his heels in one of two ways: he might be tempted to conclude either that we do not talk about the same topic, or that his opponent has a poor grasp of the relevant concept and that the opponent's intuitions therefore can safely be ignored. Less dogmatically, he could move on and provide what we claim is required: further argument. And that would be wise, for there is no safety in dogmatism here.

Imagine that I say: "To say something is *poisonous* normally implies it is not wise to eat it. However, there are uses of *poisonous* without such a consequence. So an opponent of my view might feel safe in just digging his heels ... but it would be wise to provide an argument, for there is no safety in dogmatism here." Would it not be wiser for me to provide an argument for *my* claim in the first place?

But be this as it may, let us see whether there is an argument that a classification of utterances into semantically correct and semantically incorrect has normative consequences. What does it take for a classification to have normative consequences, to imply an *ought*? Consider the claim that killing is wrong. What argument can we bring for *it* having normative consequences? Aside of an appeal to intuition (which is usually not very helpful) we can point out the fact that if a society accepts this norm, then when a member of the society *does* kill somebody, he will suffer exclusion from the society (he will not only be despised, but excluded literally, jailed, if not executed); and if not immediately, then surely upon recidivating he will be excluded from the society permanently. What does it mean that the claim that moving a rook diagonally is wrong has normative consequences? Well, whoever does move a rook in this way will be excluded from the present chess game, and upon recidivating will, sooner or later, bring about her exclusion from the community of serious chess players permanently. And in a very similar way, whoever uses English words in a wrong way repeatedly will be excluded from the community of English speakers—even if the noises he emits still sound like English words, they will not be taken seriously as English pronouncements.

Obviously these cases differ with respect to the kind of sanctions involved. There are norms whose violation may cost you your life and there are those that you can violate with only mild consequences. But in all cases you risk exclusion from a circle of adherents of the corresponding norm—be it a whole society or a smaller circle, such as that of chess players or of English-speakers. Hence the claim of Glüer and Wikforss that only some of these cases concern genuine, prescriptive correctness lacks clear substantiation.

Hattiangadi (2006) offers what looks like a more elaborated argument for the same thesis. She writes (238):

For a rule to be prescriptive, it must tell me what I ought to do. According to MP [the assumption that if *t* means *F* then it applies correctly to something iff it is an *F*], the meaning constituting rule for 'horse' must imply that 'horse' applies correctly to all and only horses. However, it is not the case that I ought to apply 'horse' to all and only horses—I am not obligated to apply 'horse' to all horses because I cannot do so, and 'ought' implies 'can'. The weaker rule,

stating that I should apply ‘horse’ only to horses cannot constitute the meaning of ‘horse’. The rule that tells me to apply ‘horse’ only to horses does not distinguish between my meaning horse by ‘horse’ and something else, such as brown horse or black horse. Moreover, I sometimes ought to tell lies, or use my words incorrectly to prove a point or make people laugh, and the fact that I ought to use my words incorrectly does not imply that I do not mean what I ordinarily mean by my words. Since the weaker rule cannot distinguish my meaning any number of things by ‘horse’, it simply cannot constitute meaning. Thus, Prescriptivity must be false, and meaning is not normative in the sense that is required to generate a presumption against naturalism.

I will forego discussing the claim that one sometimes “ought to tell lies, or use words incorrectly”—for this is an objection we have already encountered in Boghossian’s paper, with the result that this simply and unproblematically follows from the fact that rules can be overridden by other rules.

As follows from the discussion in the previous section, the first sentence of this argument is true only if “telling me what I ought to do” encompasses also “telling me what I ought *not* to do”. And it is clear that if English is to involve a rule describable as “‘horse’ applies correctly to all and only horses”, this cannot be construed as stating that I am to say “(lo, a) horse (!)” whenever I am confronted with a horse. Rather such a rule would tell us that pointing at a horse is incompatible with denying “This is a horse”—i.e. that if we do the former, we ought to avoid the latter.

Hattiangadi argues that though the following principle is basically valid (where “*t* is a term, *F* it’s meaning, and *f* the feature or collection of features in virtue of which *F* applies”):

*Correctness* :  $S$  means  $F$  by  $t \rightarrow (x)(S \text{ applies } t \text{ 'correctly' to } x \leftrightarrow x \text{ is } f)$ ,

its prescriptive versions are invalid:

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

*Prescriptivity\** :  $S$  means  $F$  by  $t \rightarrow S \text{ ought : } (x)(S \text{ applies } t \text{ to } x \leftrightarrow x \text{ is } f)$ .

Hence she drives a wedge between correctness and prescriptivity, which lets her claim that though we can speak about correctness in connection with meaning, we cannot speak about prescriptivity, and hence nor about normativity of meaning.

But I have already explained why I do not think we can disconnect correctness from prescriptivity.<sup>11</sup> Hattiangadi argues that moving the *Correctness* condition to the *Prescriptivity*(\*) one, we lose soundness. But I find this questionable. If we are to accept her objection to *Prescriptivity*(\*), then I think we would be able to formulate a parallel objection to *Correctness*, and hence it would seem that soundness is *not*

<sup>11</sup> Unlike Glüer and Pagin, Hattiangadi does give an example of the employment of the word “correct” which, according to her, does not imply any prescription. She invites us to imagine a theme park where there is a minimum height requirement for some of the more dangerous rides; and she says that to say that somebody is of the *correct* height in this respect does not imply any prescription. This seems to me to be utterly odd, for surely saying that somebody is of the correct height is tantamount to saying that she is permitted to enjoy the ride (i.e. that if she is *not* of the correct height, she *ought not* to enjoy it).

lost along the way from *Correctness* to *Prescriptivity*. To see this, let us clarify what exactly *Correctness* says in the first place. To do this we must get clear about the meaning of the phrase “*S* applies *t* ‘correctly’ to *x*” which it contains.

As Hattiangadi’s *Correctness* is supposed to be logically regimented, it would seem that we should take it at face value: *applying ‘correctly’* is a ternary relationship, a subspecies of *applying*. However, this would imply that whoever applies something ‘correctly’, applies it, and this would in turn imply that *Correctness* could not be true unless whoever means *F* by *t* applies *t* to every *f*; which is the objection Hattiangadi raises against *Prescriptivity*, so as she wants to accept *Correctness* and reject *Prescriptivity*, this cannot be the reading she has in mind.

It seems that any other reading would have to ascribe “*S* applies *t* ‘correctly’ to *x*” a more complicated logical structure. But as analyzing it as “if *S* applies *t* to *x*, then this is correct” or “*S* applies *t* to *x* and this is correct” would clearly not do, there seems to be only one viable candidate—the counterfactual one: “if *S* were to apply *t* to *x*, then this would be correct”. But once we read *Correctness* in this way, there would seem to be a parallel articulation of *Prescriptivity*; we just have to acclimatize to the fact that since the antecedent of *Correctness* restricts it to situations where the application is in place, its consequent cannot talk about an obligation to apply it. It can at most state in which cases the application is incorrect—i.e. when it ought *not* to have been in place. Hence we have

$$\begin{aligned} \text{Correctness}^{**} : S \text{ means } F \text{ by } t &\rightarrow \\ &(x)(\text{if } S \text{ were to apply } t \text{ to } x, \text{ then this would be correct} \leftrightarrow x \text{ is } f), \\ \text{Prescriptivity}^{**} : S \text{ means } F \text{ by } t &\rightarrow \\ &(x)(\text{if } S \text{ were to apply } t \text{ to } x, \text{ then } S \text{ ought not to have done so} \leftrightarrow x \text{ is not } f), \end{aligned}$$

which can be reformulated in a more natural way<sup>12</sup>:

$$\begin{aligned} \text{Correctness}^{***} : \text{If } S \text{ means } F \text{ by } t, \text{ then it is correct for } S \text{ to apply } t \text{ to } x \text{ if, and} \\ &\text{only if, } x \text{ is } f. \\ \text{Prescriptivity}^{***} : \text{If } S \text{ means } F \text{ by } t, \text{ then } S \text{ ought not to apply } t \text{ to } x \text{ if, and} \\ &\text{only if, } x \text{ is not } f. \end{aligned}$$

The latter then can be formulated also in the ‘permissive mode’, as<sup>13</sup>

$$\text{Prescriptivity}^{****} : \text{If } S \text{ means } F \text{ by } t, \text{ then } S \text{ may apply } t \text{ to } x \text{ if, and only if, } x \text{ is } f.$$

Hence it would seem that this drives out any putative wedge from between correctness and prescriptivity, which might have been produced by the above compromised regimentation.

<sup>12</sup> Though the “if” direction of the biconditionals is straightforward, the “only if” might be felt as less perspicuous. But this direction is not important and in fact the “and only if” can be deleted—if we know everything that is forbidden, what is not forbidden follows. Hattiangadi’s (206, 238) objection that “the rule that tells me to apply ‘horse’ only to horses does not distinguish between my meaning horse by ‘horse’ and something else, such as brown horse or black horse” would be relevant only if this were the only rule governing *horse*. But the inferential pattern governing *horse* and thus constituting its meaning involves the rules that we can infer (*This is a*) *horse* from (*This is a*) *brown horse*, but not vice versa.

<sup>13</sup> In this way, we have reached, via a different route, the same result as Whiting (2009).

In this way, the fact that *t* means *F* does not imply a command, it does not tell what one *ought* to do; but it does imply a normative, for it tells us what one *ought not* to do; or, for that matter, what one *may* do. And unless we mistake commanding for general normativity, we can see that *Correctness* and *Prescriptivity* are two sides of the same coin.

Hence once again, I do not disagree with Hattiangadi insofar as she claims that there are no rules of language which function as commands guiding our linguistic activities. However, the trek from the claim about linguistic rules not being commands, which I take to be established, to the general thesis that *meaning is not normative* is long and, I am convinced, unassailable. And I do not see that Hattiangadi has even set out on it in her paper.

### What Do the Normatives ‘Really’ Say?

Let us consider an objection Hattiangadi (2003) launches specifically against Brandom’s concept of normative attitudes:

[I]t is unclear how Brandom’s view differs from a straightforwardly naturalistic one. The starting point is supposed to be a proto-hominid community in which there are norms, but no concepts or contents—i.e., neither propositional attitudes, nor explicit thoughts. Brandom says, ‘the account of norm-instituting social practices must appeal to capacities that are plausibly available in primitive prelinguistic cases, and yet provide raw materials adequate for the specification of sophisticated linguistic practices, including logical ones.’ The key, according to Brandom, is to look at ‘assessments of propriety’, at ‘attitudes of taking or treating performances *as* correct or incorrect’. And although Brandom uses normative vocabulary to say that the proto-hominids treat each other’s performances as ‘correct’ or ‘incorrect’, he suggests that they do so by way of their purely *physical* behaviour and abilities.

The last sentence is odd: it suggests that besides “*physical* behavior and abilities” there would exist some other (*nonphysical?* *metaphysical?*) kind of behavior. However, the claim of the inferentialist is not that there would be two kinds of behavior, one physical and one not, but rather that there may be certain very complex patterns of behavior that are not usefully addressable in other than normative terms, which we are not able to translate into non-normative vocabulary. (While from the ‘outsider’ perspective, the normative vocabulary ‘merely’ proves itself to be a handy—if not a practically indispensable—tool of description, from the ‘insider’ perspective it is the essential means of carrying out the speech act into which describing, contaminated by endorsing, mutates.)

Hattiangadi continues that “it appears as though Brandom is offering a dispositionalist account of the determination of correctness—since the starting point includes nothing more than behavioral dispositions”, based merely on “the capacity for responsive discrimination”. It is true that from the ‘outsider perspective’ we may be able to describe the practices of a community of norm-endorsing individuals in non-normative terms. After all, it is clear that an external observer can observe nothing else than (very complex) behavioral patterns.

But we should pay attention to what does *not* follow from this. Firstly, it does *not* follow that the observer from the ‘outsider perspective’ would be able to make do with a ‘naturalistic’ vocabulary in some narrow sense of the word (like the vocabulary of physics). I think that an analogue of Davidson’s (1974, p.154) argument to the effect that a ‘vocabulary of agency’ (Ramberg’s 2000, term) is irreducible to naturalistic vocabulary because it serves a different aim, applies here. (Even somebody with no idea what chess is can come to understand and explain what chess players do—but can they do so without employing such terms as *rules*, *error* etc.—i.e. using only a straightforwardly naturalistic vocabulary?) However, this is not the most crucial problem.

Secondly and more importantly, it does *not* follow that we can translate true normatives into a non-normative idiom. When I claim that killing this cat would be wrong, I am not reporting a behavioral pattern which I instantiate. Instead, what I am doing—perhaps among other things—is endorsing a rule, and endorsing is not describing. True, an observer from outside of my community may report on this: *Peregrin endorses a rule according to which killing the cat he is pointing at would be wrong*, which could perhaps be further rendered as *Peregrin instantiates such and such kind of behavioral pattern centered around killing animals (involving not killing them, diverting others from doing so, ostracizing those who do kill one, but involving also complicated adaptive feedback loops ...)*, and now he emits a sound aimed at making his peers display the same pattern and consequently also not kill the cat he is pointing at, but nothing like this yields us a *translation* of my utterance.

Thirdly and most importantly, I can assume an ‘outsider perspective’ w.r.t. a set of rules and report on them from this perspective only via acquiescing in some other set of rules. (I can report on a chess game without endorsing the rules of chess; distancing myself from the rules of chess would merely make it impossible for me to check one of the players or to castle. But the very practice of *reporting* presupposes a language game, that is, according to the inferentialist, essentially rule-governed.) And even if we waived this problem, the inferentialist’s conviction is that we live in a world where many of the inhabitants (other people as *thinking, rational and responsible persons*; their antics as *intentional deeds*; the sounds they emit as *meaningful utterances*; ...) are constructed in terms of our normative attitudes and as a result are, in Sellars’s often quoted words “fraught with ought”. Trying to describe them with purely naturalistic vocabulary condemns us to describing a drastically impoverished *simulacrum* of our world. (This has to do with the dialectics of what Sellars 1962, called the *manifest image* and the *scientific image* of the world.)

The upshot is that even if we were able to give a naturalistic account of some kind of normative attitudes (from the ‘outside perspective’: a possibility I do not want to dismiss though I have tried to indicate that it is far from straightforward), we could do so only by falling back on other normative attitudes—those that keep us in the business of *giving accounts*, and indeed living within our human world. As Lance (2000, 134) puts it, “to *talk about* rules is to move *outside* the talked-about rules *into* another framework of living rules.”

Let us return to Hattiangadi. She continues (*ibid.*, 426):

Even though the dispositions of members of a community are described in normative terms, the normative vocabulary affords no purchase on the

problem. The reason is that the one thing Brandom *cannot* assume is that the pre-conceptual abilities of creatures includes the ability to *think*, i.e., to entertain concepts. This is not just because they are supposed to have only pre-conceptual abilities, but also because, if it were necessary that they *think*, Brandom would face the problem of regress. Moreover, the leading idea of this work is that the rules that determine conceptual content are *implicit* in practices. The account of how a practice must be in order for it to institute conceptual content cannot presuppose that the participants of the practice can have explicit, contentful thoughts.

This indicates that either Brandom's elucidation is circular (it presupposes the ability to think which is among the things it is to explain), or it is idle. The same challenge is articulated more concisely by Glüer and Wikforss (2009, 57), who portray the normativist as facing the "dilemma of regress and idleness":

[T]he regress of contents forces us either to give up CD normativism [a doctrine "according to which content is determined by norms in the first place"—J.P.] or to accept that there is rule-following that does not require any kind of prior intentional state. There are no other options. And given how overwhelmingly intuitive the idea of an intentional condition on rule-following is, for the determined CD normativist, 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 condition goes by the board?

It is quite clear that the inferentialist can identify rule-following neither with mere regularity nor in general with an intentional following of an explicit rule—if these were the only options, then a vicious circle would indeed be inevitable. Hence the inferentialist's stance indeed stands and falls with the existence of something in between the Scylla and the Charybdis—something like Sellars's "patterned-governed behavior" (which, as he puts it, is neither "merely rule conforming", nor "fully rule obeying").

There is not enough room to explain the nature of this kind of behavior here (I have done so elsewhere—see Peregrin 2010a) For the sake of the present discussion, just consider the musings of Wittgenstein's leading him to the conclusion that there must be rules that I follow *blindly*. Suppose we cannot follow a rule unless we assume an intentional attitude to it. What exactly is it that the intentional attitude is directed at? In the simplest case it is a normative, an instruction in a natural language; and in any case it must be some token representing, expressing or envisaging the rule. Now any such token, to lead us to the rule, must be *interpreted*, and to lead us to the rule in question, it must be interpreted *in a particular way*. Hence, taking the attitude to the token as the attitude directly to the rule presupposes that we interpret the token in a *correct* way. In other words, either we need a rule to follow a rule, which leads us into a regress, or we can have the ability of doing correct things without thereby following the rule.<sup>14</sup>

<sup>14</sup> As Wheeler (1986, 492) points out (in a different context), "speech and thought are brain-writing, some kind of tokenings which are as much subject to interpretation as any other".



From this viewpoint, the first horn of this dilemma is clearly unacceptable—it would render rule following utterly impossible, whereas there is hardly a way to deny that we do follow many rules. Hence we must embrace the second. There must be some way of setting rules into motion; some kind of rule following that is not yet conscious obeying of explicit rules.

### From Regularities to Rules

Suppose that members of a tribe start to use a certain kind of sound when pointing at something edible, and that they start to deter anybody from using the sound in a different situation (say by beating her with sticks). Of course, at this point there is as yet no genuine rule, but merely a regularity, i.e. something describable exhaustively in the behaviorist idiom and incapable of conferring anything like a real content (i.e. content of the kind of that which the words of our languages have) on the sound. (And if we decide to call the acts of diverting members of the tribe from using the sound in certain ways ‘punishments’, then it is nothing more than a metaphor.)

What may happen then, however, is that what starts as being a *rule* or a *punishment* only metaphorically, develops into something that *is* a rule or a punishment, and hence into something that is, as such, describable only in terms of a normative vocabulary. How does this transition happen? Well, it is a kind of a bootstrapping process involving the development of language, during which the participants are gaining the ability to classify situations in certain ways and thereby articulate and grasp, to use Sellars’s (1969) phrase, various *ought-to-be*’s, while at the same time gaining the ability to engage within genuinely normative practices. (The two abilities are inextricable, for to be able to articulate situations we need meaningful language, which is a matter of advanced normative practices, whereas to have such advanced normative practices, we need language. However, bootstrapping is not circularity.) Members exhibiting truly *normative* attitudes not only divert others from doing certain things, but do it because they see it as violation of rules that, according to them, should be accepted.

Consider the process whereby some originally unclear and only rudimentary rules of an imitation football game played by a group of kids in a backyard may develop into wholly explicit and unambiguous rules. The process is that in which the original fuzzy and negotiable rudiments of rules first crystallize into relatively clear and respected, though still unwritten, rules; which then become explicit in that the players come to articulate them explicitly and may even have them written down. What counts as clear rules are not only the explicitly articulated ones, but also their clear and generally respected, though still unwritten predecessors.

Let me point out that describing the just presented transition faces a certain “conceptual problem”, clearly diagnosed by Davidson (1999)<sup>15</sup>—the development from what only *looks like* a rule (i.e. forms of behavior of non-normative, non-minded and non-conceptual creatures) to what *genuinely is* a rule (i.e. actions of minded and conceptual creatures capable of binding themselves with rules) takes

<sup>15</sup> Davidson, of course, does not talk about rules; however, the predicament of his ‘vocabulary of agency’ is, from this viewpoint, the same as that of our normative vocabulary.

place within the no-man's land between the physicalistic realm of law and the normative realm of freedom, a land which is not describable in terms of either of its neighbours. As Davidson characterized the situation:

The difficulty in describing the emergence of mental phenomena is a conceptual problem: it is the difficulty of describing the early stages in the maturing of reason, the stages that precede the situation in which concepts like intension, belief and desire have clear application. ... What we lack is a satisfactory vocabulary for describing the intermediary steps. ... We have many vocabularies for describing nature when we regard it as mindless, and we have a mentalistic vocabulary for describing thought and intentional action; what we lack is a way of describing what is in between. (ibid., p. 127–8)

As I have put it elsewhere (see Peregrin 2006), the bridge from the non-normative to the normative can never really be built: for if we try to do so using physicalist terms, we cannot reach the normative shore; whereas using normative terms we are there before we start.

Hence our predicament in this respect is, I think, the very one which was so vividly pointed out by Davidson: interpreting the behavior of other people is an enterprise quite different from the enterprise of natural science, which aims to subsume the episodes of the objects of nature under strict causal laws. Again, a similar point was made by Sellars (1962), who urged the irreducibility of the manifest image, in which there is room for normativity and hence such things as meanings or persons, to the scientific image, which maps the filaments of causal relationships between things. (Also we should not overlook Kant (1790, §II), with his distinction between the realm of the concept of nature and the realm of the concept of freedom, and his insistence that “the concept of freedom just as little disturbs the legislation of nature, as the concept of nature influences legislation through the concept of freedom”.)

Are normative attitudes simply cases of “responsive discrimination”? Of course they *presuppose* some abilities of responsive discrimination, and indeed they *can* be seen as cases of such discrimination. However, and this is crucial, they are more than this. We have already noticed that the discrimination between the *correct* and the *incorrect* differs from a more ordinary discrimination (such as the discrimination of the difficult from the easy): in the latter case, the criterion of the discrimination is independent of the person who does the discrimination and of the process of discrimination, whereas in the former one the process takes part in constituting and upholding the criterion. We have stressed that by judging something to be correct or incorrect we not only state the fact that it is such, we also endorse the rule which underlies such a verdict. By every act of this normative kind of classification we not only classify, but sustain—or, as the case may be, amend—the classificatory criterion.

### **The Normativity of Meaning Claimed by Inferentialism**

The discussions of the objections to the normativity of meaning in the previous sections have revealed the notion of normativity of meaning (and normativity in general) that is, I believe, inherent to the inferentialist standpoint. The notion we

have reached is close to the one exposed by Lance and O’Leary-Hawthorne (1997). Like them, I hold for obvious that normativity is a matter of rules established by us, humans—hence that we must reject what they call the “transcendental conception” of normativity, i.e. the view according to which ‘what is correct’ is essentially independent of what we people really do. Similarly, I concur with their rejection of what they call the “attributive conception”, where ‘what is correct’ always boils down to what is legitimate according to the current rules of such or another community.<sup>16</sup>

Lance and O’Leary-Hawthorne (*ibid.*, 202–3) explain the peculiar status of the true normatives so that they are somewhere in between indicatives and imperatives:

Normatives are in many ways just like ordinary declaratives. They take their place in the game of giving and asking for reasons, serving as premises and conclusions in reasoning. ... But unlike declaratives, one of the direct, and widely stable, consequences of application of a normative is the appropriateness of some act; to commit oneself to a normative is *ipso facto* to commit oneself to the propriety of some act.

Hence, just like indicatives, normatives can be *justified* and are thus subject to the game of giving and asking for reasons. (Imperatives are different; they may be ‘justified’ only in the sense that their utterer possesses a certain social status that grants him the power or the right to give orders or instructions to his audience.) Like imperatives, on the other hand, they establish commitments and/or entitlements (over and above the standard commitments established by indicatives, *viz.* the commitment to provide reasons if challenged and the entitlement to reassert it deferring a justification to the original asserter).

The crucial feature of these speech acts that provides for their peculiarity and which provides for the middle course between the “transcendental” and the “attributive” conception is the fact that rules in human communities are always in a sense open. They are both open to revisions and also open insofar as a rule can never settle everything w.r.t. the kind of behavior that is its target. Both the range of cases it applies to and the way it applies to them are subject to negotiation—and a true normative, which involves an application of a rule, always either reinforces the rule as it stands, or provides for its modification, extension or specification.

However, this must not be read in the sense that an utterer of a true normative is free to create norms quite deliberately. First, normatives must be anchored in the existing practices with existing rules and though they may, and usually do, go beyond them, they can do so only to such an extent that it makes sense to say that they are exercisings of the existing rules. Second, if a normative aims at a modification or an extension of the existing practices, it counts as a proposal, which

<sup>16</sup> In fact I think (and here I may differ from Lance and O’Leary-Hawthorne) that some utterances of normatives *should* be read “attributively”—namely those that are done from the ‘outsider perspective’. Hence to say that an expression of a language is used correctly thus and so may simply mean a report on the fact that this use is legitimate according to the correct rules of the relevant linguistic community. In this context, one may perhaps dispute whether the normative vocabulary is truly crucial. What is, however, crucial is that the same normative may also be used to carry out the peculiar kind of speech act in which a declarative ingredient is coupled with an endorsement.

can be taken as established only if it survives any occurring criticism and if it comes to be generally accepted.<sup>17</sup>

Moreover, though we can say that this stepwise development of rules amounts to a *creation*, it is usually understood as a case of *discovery*. To explain this peculiar feature of rules, compare their status with the status of certain objects of mathematics. On the one hand, it is acceptable to say that it was Cantor who devised sets, or that it was Galois & comp. who invented groups. But on the other hand, these mathematical objects were devised as entities which exist timelessly, and as such *cannot* have been *brought* into being. Thus, once we devise them, we must look at them as having been *discovered* by us, while having been here all along. And in a similar way, as correctness is something essentially persistent, once we extend or modify a rule we look at the kind of correctness instituted by it as something that has been here all along and that we should use to evaluate also past actions. Therefore Lance (2000, p. 127) can say: “The rule governing a use of a bit of language does not have to be *there* as a social regularity, or an ideal in anyone’s mind prior to justifying the claim that a person was already committed to it all along.”

The fact that rules are not always explicit is sometimes expressed, as we have also done, by saying that they are a matter of *practices*. The important thing about practices is that they institute ‘tracks’ of behavior a deviation from which counts as an error. And, of course, practices need not only be abstract or linguistic (like adding and subtracting numbers); they may essentially involve the extralinguistic world. As Rietveld (2008) puts it, they may be “situated” (while at the same time essentially normative):

Given that normativity is constituted by the communal custom in which the individual’s performance is embedded, mistakes by the skilled individual are possible. Thus, even though in unreflective action a performance is not undertaken for any explicit reasons, it can fail. It is thanks to this complex context that the skilled individual’s unreflective performance, which typically is, in a sense, nothing but a ‘blind’ response to relevant affordances (namely the individual’s being moved to respond by them), is normative nevertheless.

Glüer and Wikforss (2009, 58) think otherwise:

A Practice, after all, is a regularity in behaviour (social or individual) and this notion cannot be employed to secure the distinction between merely acting in accordance with a rule, acting in regular ways, and being guided by a rule.

Well, admittedly, we can say that a practice *can be seen* as a regularity in behavior. But this does not yet mean that this concept cannot help us make the differences we claim: we must distinguish between the regularity of the very behavior regulated by the rules (e.g. the drivers driving mostly on the right side of the road w.r.t. the corresponding rule of traffic) and regularities of a much more exclusive behavior centered around this one, i.e. additional regularities in the behavior of those who react to the behavior that is not regular in the above way, regularities in behavior of those who react to these reactions etc. (e.g. the behavior of policemen who fine those who do not drive on the right side, journalists who

<sup>17</sup> Cf. Gauker’s (2007) concept of *circle of deference*.

criticize the policemen for being over- or under-reactive etc.) Moreover, the behavior patterns will display complicated feedback loops (corresponding to the fact that a rule is always ‘in the making’ ...). This is what makes a practice *such a complicated* behavioral pattern (a motley of patterns?) that any attempt to capture it in a naturalistic idiom (even from the ‘outsider’ perspective) seems futile.

Moreover, we may be content with saying that a practice is a regularity of behavior if we are content with staying disengaged. But what we call *practices* are often behavioral patterns to which we do not intend, and sometimes even are not able, to stay disengaged. Though what we do when expressing them in words may look like giving a description, what we really do involves, due to their essential ‘openness’, extrapolative proposals.

Rouse (2007, p. 49) characterizes the relationship between practices and normativity in the following illuminating way:

A normative conception of practices makes normativity irreducible but not inexplicable. There are at least three crucial aspects to its explication of normativity. First, the bounds of a practice are identified by the ways in which its constitutive performances bear on one another, rather than by any regularities of behavior or meaning that they encompass. One performance responds to another, for example, by correcting it, drawing inferences from it, translating it, rewarding or punishing its performer, trying to do the same thing in different circumstances, mimicking it, circumventing its effects, and so on. ... A second crucial feature of practices, normatively conceived, is that these patterns of interaction constitute something at issue and at stake in their outcome. ... Normative practice theories, however, take the issues and stakes in practices to be not merely subject to epistemic uncertainty, but perspectivally variant or opentextured, and this amounts to a third crucial feature of their conception of practices.”

Rouse also invokes the terms “mutual accountability” and “diffraction” for the way in which the components of the practices bear on each other, producing the normative effects. Practices involving the normative attitudes and thus instituting correctness are indeed “diffractive”—they are a matter of interaction between moves, counter-moves, counter-counter-moves etc. such that counter-moves not only push themselves off the moves to which they react, but at the same time push these moves to make the whole system interdependent in all kinds of dimensions (somewhat like billiard balls ...).

### Conclusion: Why Meaning is Normative

Let me now summarize the sense in which meaning, from my inferentialist perspective, is normative:

1. *It is rules (especially inferential rules<sup>18</sup>) that constitute semantics.* They constitute it in a similar sense to how the rules of chess constitute *kings*, *pawns* and *bishops*—by instituting certain roles for tokens governed by them. Meaning claims are claims to the effect that a token is of a type, that it plays a specific kind of role.

<sup>18</sup> Note that nothing in this paper hinges on *inferential* rules being crucial. That is why I say that the argumentation presented here does not presuppose an inferentialist standpoint in a narrow sense.

2. *The fact that the rules constitute meanings does not rob them of their normativity—they are constraints that render certain actions as incorrect (though on a level that recedes to the background of meaningful talk).* We must not be misled by the fact that if we decide to move ‘inside’ the ‘space’ constituted by the rules, they move out of our sight.
3. *Normativity is ultimately grounded in normative attitudes; which, however, does not allow us to reduce the normative to the non-normative, because of the ‘diffractive’ nature of the attitudes.* Reducing normativity to the attitudes thus help us merely reduce complex forms of normativity to more elementary forms.<sup>19</sup> The impossibility of the reduction to the non-normative is not a matter of instigating two incommensurable strata of reality, but rather that the normative and the non-normative idioms constitute two different kinds of speech acts, and that the enterprise of accounting for what is correct, for commitments and entitlements, or for persons being responsible to each other, is incommensurably distinct from accounting for causal relationships.
4. *The ultimate effect of the rules consists in their constrictive import opening up a brand new space for actions not previously available.* Rules institute roles; and the actions that are regulated by the rules essentially involve the role-bearing items as such; they cannot be seen as some old actions in new guises.

The fact that the entities populating the space thus opened up by the system of rules, and being the proper subjects of the rules, are themselves *constituted* by the rules does not mean that they are unreal and can only parasite on something that is real. This is analyzed in detail by Haugeland (1998): constitution, as he puts it, must not be seen as a creation *ex nihilo*, but rather as ‘letting be’—as an acknowledgment of something that, despite not being able to be what it is without us, is nevertheless independent of us.

These considerations may sound rather esoteric; but I think that chess provides us with a vivid illustration. The rules of chess that make up the game obviously open up a space of new actions and enable the players to engage in exciting activities not previously available to them. And the inferentialist idea is that, analogously, the rules of language open up the space of meaningfulness that enables us to engage in practices not previously available to us—practices of distinctively human communication, rational deliberation, building theories etc.

Hence meaning, for the inferentialist, is normative because any claim to the effect that an expression means whatever it does mean is a *normative* claim; it tells us how the expression is *correctly* used, or which *rules* it is governed by. From this viewpoint even the very claim that meaning is normative may be misleading: semantic claims do not talk about ‘normative entities’ attached to expressions, they prescribe how to handle the expression in the proper way.

## References

- Boghossian, P. A. (1993). Does an inferential role semantics rest upon a mistake? In A. Villanueva (Ed.), *Philosophical issues* 3 (pp. 73–88). Atascadero: Ridgeview.

<sup>19</sup> In this respect it is a reduction similar to Davidson's reduction of his vocabulary of agency to the relationship of *holding true*.

- Boghossian, P. A. (2005). Is meaning normative? In C. Nimtz & A. Beckermann (Eds.), *Philosophie und/als Wissenschaft* (pp. 205–218). Paderborn: Mentis.
- Brandom, R. (1994). *Making it explicit*. Cambridge: Harvard University Press.
- Chisholm, R. M., & Sellars, W. (1958). Intentionality and the mental: Chisholm-sellars correspondence on intentionality. In H. Feigl, M. Scriven, & G. Maxwell (Eds.), *Minnesota studies in the philosophy of science* (Vol. II, pp. 521–539). Minneapolis: University of Minnesota Press.
- Davidson, D. (1974). Belief and the basis of meaning. *Synthese* 27, 309–323. Reprinted in and quoted from Davidson: *Inquiries into truth and interpretation* (pp. 141–155). Oxford: Clarendon Press, 1984.
- Davidson, D. (1999). The emergence of thought. *Erkenntnis*, 51, 7–17. reprinted in and quoted from Davidson: *Subjective, Intersubjective, Objective*, Clarendon Press, Oxford, 123–134.
- Fodor, J., & LePore, E. (1993). Why meaning (probably) isn't conceptual role. In E. Villaneuva (Ed.), *Science and knowledge* (pp. 15–35). Atascadero: Ridgeview.
- Gauker, C. (2007). The circle of deference proves the normativity semantics. *Rivista di Estetica*, 47, 181–198.
- Glüer, K., & Pagin, P. (1999). Rules of meaning and practical reasoning. *Synthese*, 117, 207–227.
- Glüer, K., & Wikfors, A. (2009). Against content normativity. *Mind*, 118, 31–70.
- Hattiangadi, A. (2003). Making it implicit: brandom on rule following. *Philosophy and Phenomenological Research*, 66, 419–431.
- Hattiangadi, A. (2006). Is meaning normative? *Mind & Language*, 21, 220–240.
- Hattiangadi, A. (2009). Some more thoughts on semantic oughts: a reply to Daniel Whiting. *Analysis*, 69 (2009), 54–63.
- Haugeland, J. (1998). Truth and rule following. In: *Having thought*. Cambridge: Harvard University Press.
- Kant (1790). *Kritik der Urteilskraft*, Berlin.
- Lance, M. (2000). The word made flesh: toward a neo-sellarsian view of concepts, analysis, and understanding. *Acta Analytica*, 15, 117–135.
- Lance, M., & O'Leary-Hawthorne, J. (1997). *The grammar of meaning*. Cambridge: Cambridge University Press.
- Peregrin, J. (2006). Brandom and Davidson: what do we need to account for thinking and agency. *Philosophica*, 75, 43–59. <http://logica.rug.ac.be/philosophica/>.
- Peregrin, J. (2009). Inferentialism and the compositionality of meaning. *International Review of Pragmatics*, 1, 154–181.
- Peregrin, J. (2010a). The enigma of rules. *International Journal of Philosophical Studies*, 18, 377–394.
- Peregrin, J. (2010b). Semantics without meaning? to appear. In R. Schantz (Ed.), *Prospects of meaning*. Berlin: de Gruyter.
- Ramberg, B. (2000). Post-ontological philosophy of mind: Rorty versus Davidson. In R. Brandom (Ed.), *Reading Rorty*. Blackwell: Oxford.
- Rietveld, T. (2008). Situated normativity: the normative aspect of embodied cognition in unreflective action. *Mind*, 117, 973–1001.
- Rouse, J. (2007). Social practices and normativity. *Philosophy of the Social Sciences*, 37, 46–56.
- Searle, J. R. (1969). *Speech acts: an essay in the philosophy of language*. Cambridge: Cambridge University Press.
- Sellars, W. (1949). Language, rules and behavior. In S. Hook (Ed.), *John Dewey: Philosopher of science and freedom* (pp. 289–315). New York: Dial Press. reprinted in Sellars (1980).
- Sellars, W. (1962). Philosophy and the scientific image of man. In R. Colodny (Ed.), *Frontiers of science and philosophy* (pp. 35–78). Pittsburgh: University of Pittsburgh.
- Sellars, W. (1969). Language as thought and as communication. *Philosophy and Phenomenological Research*, 29, 506–527.
- Sellars, W. (1974). Meaning as functional classification. *Synthese*, 27, 417–437.
- Sellars, W. (1992). *Science and metaphysics*. Atascadero: Ridgeview.
- Wheeler, S. C., III. (1986). Indeterminacy of french interpretation: Derrida and Davidson. In E. Lepore (Ed.), *Truth and interpretation: Perspectives on the philosophy of Donald Davidson* (pp. 477–494). Blackwell: Oxford.
- Whiting, D. (2009). Is meaning fraught with ought? *Pacific Philosophical Quarterly*, 90, 535–555.
- Wikfors, A. M. (2001). Semantic normativity. *Philosophical Studies*, 102, 203–226.
- Williams, B. (1973). Deciding to believe. In B. Williams (Ed.), *Problems of the self* (pp. 136–151). Cambridge: Cambridge University Press.
- Wittgenstein, L. (1953). *Philosophische untersuchungen*. Oxford: Blackwell.
- von Wright, G. H. (1963). *Norm and action*. New York: Humanities.