

Inferentialism and the Normativity of Meaning

Jaroslav Peregrin

peregrin@ff.cuni.cz webpage: <http://jarda.peregrin.cz>

Institute of Philosophy, Academy of Sciences of the Czech Republic, Jilská 1, 110 00 Praha 1, Czech Republic

Abstract. The paper addresses some frequent objections to the claim that meaning is normative, thus defending the inferentialist construal of meaning that does entail this claim. The objections we discuss are (i) that there is no norm that assertions should aim at the truth, (ii) that there are no norms commanding us how to speak and (iii) that a normative account is bound to collapse into a naturalistic one. We conclude that the way in which normativity is built into the inferentialist framework is not vulnerable to any of these objections. Along the way, we try to clear up some misunderstandings surrounding the status of normativity within semantics.

1 Is Meaning Normative?

Inferentialist view of meaning (based on the assumptions 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 should be used so and so¹) faces two kinds of objections. First, there are general objections to any normative construal of meaning, and then there are more specific objections targeted specifically at inferentialism. In this paper I will address the objections of the former kind (I address some of those of the latter kind elsewhere²).

To give an example of the way a general objection to the claim that meaning is normative usually goes, let me quote one of its most influential exponents, Paul Boghossian (2005, 212):

To put the matter concisely, the linguistic version of the normativity thesis, in contrast with its mentalist version, 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.

¹ See Peregrin (2008) for a more detailed exposition of this aspect of inferentialism.

² See Peregrin (2001); and especially Peregrin (t.a.) (All my papers as well as a longer draft of this paper are available from my homepage at jarda.peregrin.cz/mybibl.php.)

Sellars (1992, p. 101), whose lead I follow here, thinks otherwise: he understands truth as correct assertability, which entails that an assertion is *correct* if what is asserted is true; and hence, in this sense, that we *ought to* assert the truth³.

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. True, on one of its readings, this may still be read as stating a kind of fact, namely that an activity within a community is - as a matter of fact - governed by certain rules; 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,

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)

I am fond of describing the situation in terms of an 'inner space' that some systems of rules have the ability to constitute. From outside of the space we can only report the fact that such or other rules are in force for the insiders; but once we join the insiders, the rules 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. Hence, let me call the former reading the 'outsider' reading and the latter the 'insider' one.

Thus, claims to the effect that something is correct or that something ought to be done (I will call them *normatives*, for short), on the insider reading, are something different from assertions or reports. They do not report that something is the case, they point out that something *ought to be* the case; hence they always involve the utterer's taking a rule for 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 (like 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 of a rule.

This, of course, is not the only thing that the 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 am claiming 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, even on their 'outsider' read-

³ I also think that there *cannot be* a norm that we should *believe* the truth, as it seems to me obvious that one is not free to decide what to believe. 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.

ing, report a specific kind of facts, *viz.* normative facts - but the possibility of taking this characterization literally is compromised by the fact that the alleged fact is a fact only as a result of a rule that does not exist quite independently of the statement, for the statement takes part in its constitution.

Hence, compare the following two claims:

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

Both claims 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 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 a rather different 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.

2 Do inferential rules guide us?

Another kind of objection to a construal of meaning as a product of rules was most clearly articulated by Glüer and Pagin (1999), who argue that "rules that can be regarded as *constitutive of* meaning, are not capable of guiding speakers in the ordinary performances of speech acts." I think that here we must distinguish between two senses of "guiding": on the most straightforward sense, in which "guiding" amounts to directly telling what to do, Glüer and Pagin are undoubtedly right. But there is a different sense of "guiding", on which "guiding" involves telling what *not* to do, and in this sense rules that are constitutive of meaning *do* guide us. These rules erect limits to what we can assert if we are to use expressions as meaning what they do⁴.

However, the rules in question come to be buried under the visible surface of our discursive practices. If we see the practices as employing expressions with specific meanings, then the rules become invisible for their being in place is part and parcel of the very meaningfulness of the discourse. Taking the expressions as having their meanings necessarily, is to take the constitutive rules as incapable of being violated, and hence as not being rules worthy of their name. However, as a matter of fact, expressions always have their meanings *contingently*, and what holds the meanings fixed is precisely our holding on the rules that constitute them, our staying within the corresponding space of meaningfulness.

Hence if Glüer and Pagin claim that "the problem with constitutive rules ... is that there simply isn't anything that they require", I see this as an illusion arising from mistaking the inside of the space of meaningfulness for a reality that was once estab-

⁴ As I think it is the actual term "rule *following*" which has given rise to the underlying confusion, I would suggest replacing it by something less confusing - perhaps "bouncing off rules"?

lished and is independent of us since. Contrary to this, I maintain that the existence of this very space is secured by the fact that the rules of language *do* require something - it is only while we accept them that the space is here in the first place.

A more radical version of Glüer's and Pagin's objection is voiced by Hattiangadi (2006). 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.

As follows from the above discussion, 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.

What, then, about Hattiangadi's claim 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"? Indeed it does not, and if this were the only rule governing 'horse', then it would indeed be indistinguishable from the more specific terms. But fortunately it is not. 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.

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. The disagreement comes when she comes to claim that *meaning is not normative*. The trek from the claim about linguistic rules not being commands, which I take to be established, to this general thesis is long and, I am convinced, unassailable. And I do not see that Hattiangadi has even set out on it in her paper.

3 What do the normatives 'really' say?

Let us now 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 atti-

tudes, 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* behavior 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 certain very complex patterns of behavior are not usefully describable in other than normative terms, which we are unable to translate into non-normative vocabulary, i.e. especially not into the vocabulary of physics. (This is to say neither that the behavior is not physical, nor that it is in no way describable in non-normative terms - it is to say merely that no its non-normative description possesses the explanatory power which we need and which we get from the normative description.)

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 can 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 'outsider' 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 (1973, p. 154) argument to the effect that a 'vocabulary of agency' (Ramberg's, 2000, term) is irreducible to the naturalistic vocabulary because it serves a different aim, applies here. (Even somebody who has no idea what chess is can come to understand and explain what chess players do - but can he do so without the employment of such terms as *rules*, *error* etc. - i.e. using only a straightforwardly naturalistic vocabulary?) However, this is a problem we can safely waive here.

Secondly and more importantly, it does *not* follow that we can translate normative idiom, as employed by the insiders, into a non-normative one. When I claim that killing this cat would be wrong, I am not reporting a behavioral pattern which I instantiate: 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 can be perhaps 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, ...)*, and now he emits a sound aiming at making his peers display the same pattern and

consequently not kill the cat he is pointing at, but nothing like this yields us a *translation* of my utterance.

And 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 myself 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, we must realize that we live in a world whose inhabitants are very much constructed in terms of our normative attitudes (other people are taken as *thinking, rational and responsible persons*; their antics, as *intentional deeds*; the sounds they emit, as *meaningful utterances*; ...) and as a result, using Sellars' often quoted phrase, we are in a world "fraught with ought". Trying to describe this with a purely naturalistic vocabulary would leave us with 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 'outsider perspective' - a possibility I do not wish to dismiss though I have tried to indicate that it is far from straightforward), we can do so only by falling back on other normative attitudes, that keep us in the business of *giving accounts*, and indeed living within our human world.

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 do more than this. We have earlier noticed how 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 making the discrimination and of the process of discrimination, whereas in the former the process itself takes part in constituting and upholding the criterion. We have stressed that by judging something to be correct or incorrect we are not only stating the fact that it is such, we are also endorsing the rules underlying such a verdict. By every act of this normative kind of classification, we are not only classifying, but also sustaining the classificatory criterion.

Suppose that members of a tribe start to use certain kinds of sounds in certain specific ways and that they start to try to stop anybody from using the sound in different ways (say by beating her with sticks). Of course, at this point we do not yet have genuine rules, but merely regularities, i.e. something describable exhaustively in the behaviorist idiom and incapable of conferring anything like a real content 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 we are only using the word as a metaphor.)

What may happen then, however, is that what can only be called a *rule* or a *punishment* metaphorically, develops into something that *is* a rule or a punishment, 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

of classifying situations in certain ways and thereby articulating and grasping, to use Sellars' (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.

4 Conclusion: why meaning is normative

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

1. It is inferential rules which constitute semantics. They constitute it similarly to how the rules of chess constitute *kings*, *pawns* and *bishops* – by instituting certain roles for tokens governed by them.

2. The fact that the rules constitute the meanings does not rob them of their normativity – they are genuine *constraints* that guide our behavior, preventing us from doing certain things. 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, but merely to reduce complex normativity to its elementary forms⁵. The impossibility of the reduction is not a matter of the fact that there would be 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 different 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, analogously, that the rules of

⁵ In this respect it is a reduction similar to Davidson's reduction of his 'vocabulary of agency' to the relationship of *holding true*.

language open up the space of meaningfulness which 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, but instead they delimit how the given expression ought properly to be handled.

References

- Boghossian, P. A. (2005): 'Is Meaning Normative?', in Ch. Nimtz & A. Beckermann (eds.): *Philosophie und/als Wissenschaft*, Mentis, Paderborn, 205-218.
- Brandom, R. (1994): *Making It Explicit*, Harvard University Press, Cambridge (Mass.).
- Chisholm, R. M. & W. Sellars (1958): 'Intentionality and the Mental: Chisholm-Sellars Correspondence on Intentionality', in H. Feigl, M. Scriven, and G. Maxwell (eds.): *Minnesota Studies in the Philosophy of Science*, vol. II, University of Minnesota Press, Minneapolis, 521-539.
- Davidson, D. (1973): 'Radical Interpretation', *Dialectica* 27, 1973; reprinted in and quoted from *Inquiries into Truth and Interpretation*, Clarendon Press, Oxford, 1984, 125-140.
- Glüer, K. & P. Pagin (1999): 'Rules of Meaning and Practical Reasoning', *Synthese* 117, 207–227.
- 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.
- Haugeland, J. (1998): 'Truth and Rule Following' in *Having Thought*, Harvard University Press, Cambridge (Mass.).
- Peregrin, J. (2001): *Meaning and Structure*, Ashgate, Aldershot.
- Peregrin, J. (2008): 'Inferentialist Approach to Semantics', *Philosophy Compass* 3, 1208-1223.
- Peregrin, J. (t.a): 'Inferentialism and the Compositionality of Meaning', to appear.
- Ramberg, B. (2000): 'Post-ontological Philosophy of Mind: Rorty versus Davidson', R. Brandom (ed.): *Reading Rorty*, Blackwell, Oxford, 351-369.
- Sellars, W. (1962): 'Philosophy and the Scientific Image of Man', in R. Colodny (ed.): *Frontiers of Science and Philosophy*, University of Pittsburgh, Pittsburgh, 35-78.
- Sellars, W. (1969): 'Language as Thought and as Communication', *Philosophy and Phenomenological Research* 29, 506-527.
- Sellars, W. (1992): *Science and Metaphysics*, Ridgeview, Atascadero.