

Normativity and the Realist Stance in Semantics

Giacomo Turbanti *
turbanti.giacomo@gmail.com

ABSTRACT

Recent attempts to define and support realism in semantics seem to acknowledge, as the only defence from skeptical attacks to the notion of meaning, a flat acceptance of the existence of representational relations between language and things in the world. In this paper I reconsider part of the mistrust about the normative character of meaning, in order to show that some of the worries urging the realists to cling on representationalism actually rest on misconceptions. To the contrary, I suggest that normativity is the main strength of a stable realist stance in semantics. Support to this suggestion comes from the reanalysis of some oft-ignored sellarsian themes.

1. Introduction

1.1. Is meaning normative?

Is meaning normative? This question has been haunting philosophy of language at least since Wittgenstein's remarks on rule-following, but, after Kripke (1982), the debate on the normativity of meaning has been hinged, to a large extent, on the problem of how to abstract linguistic rules from a naturalistic account of speakers' behavior. The frustration of this enterprise usually leads to a stark choice. The first option is to accept the normativity of meaning, but to give up on semantics: meanings are normative but we don't grasp them good enough to represent them. The second option is to keep the semantic analysis of meanings, but to try defusing underdetermination issues by separating meaning from norms. I reject both these pigeon holes, and, in what follows, I'll try to unhinge the whole framework with a sellarsian lever. Just as any lever has three parts, so too has mine: the analysis of the "Myth of

* Scuola Normale Superiore, Pisa, Italy.

the Given” will provide a fulcrum, the functional analysis of meanings will carry the load of the framework, and the analysis of normative vocabulary will allow to apply force. Once the framework will be lifted up enough, we’ll be able to see that it is nourished by roots which dig deeper into epistemology, and that its semantic consequences are just some of its more visible fruits. In fact, I claim that realism in semantics is compatible with a non reductive account of the normativity of meaning.

1.2. The skeptical stance and the realist stance

Let me try to sketch, briefly, the rough outline of the relevant part of the debate on the normativity of meaning.

The first of the aforementioned options is usually set up as a *skeptical* stance, introduced along the well-known line of the arguments in Goodman (1954), Quine (1968) and Kripke (1982). Although these arguments differ in many respects, their barebone structure is basically the same, and it goes like this: consider some linguistic content and put forward a definition for it which is *adequate* to its public usage, then either you have some independent peg to hang an analysis of such an adequacy or, inevitably, another content is deliverable which satisfies your definition while being, in fact, incompatible with the first one. Several pegs have been proposed (for instance, logical ones like “projectibility”, naturalistic ones like psychology, pragmatic ones like linguistic communities), none of which has proven to be firm enough: hence the skepticism about meaning.

These arguments seem to be compelling, so the reaction against this unwelcome result tends to be drastic. The alternative option is usually set up as a militant *realist* stance. By this side of the debate, the skeptical argument is construed, at bottom, as a semantic staging of the “naturalistic fallacy” in ethics: if meaning is an evaluative notion, then no descriptive analysis may account for it. In this sense, however, the whole charge can be dismissed simply by noticing that normativity of meaning is an unsupported assumption. The burden of the proof falls back on the skeptics who have to show where the norms come from. Indeed, the realists claim, if there are norms in the nearby of meaning, they are *constitutive* of linguistic practices, in the sense that they *describe* what it is for an expression ϕ to mean F in a given practice, e.g.,

“ ϕ ” means $F \rightarrow \forall a (\phi \text{ adequately applies to } a \leftrightarrow a \text{ is } F)$ ¹

There’s no “*ought*” hidden in such a description: here adequacy has to be measured empirically and represented extensionally.

2. Mythbusting

My purpose in this section is to provide an Archimedean point outside the conflict of these two stances. I want to make room for two claims: (a) normativity of meaning doesn’t trigger skepticism (Section 2.1); (b) normativity of meaning can be represented from a realistic perspective (Section 2.2). Together, (a) and (b) draw the outline of an argument against an insidious misconception which hides inside the formal apparatus of first order quantified logic. In my opinion, this misconception is among the main causes of contemporary relapses of the form of the Myth of the Given clearly described in §30 of Sellars (1956):

[U]nless we are careful, we can easily take for granted that the process of teaching a child to use a language is that of teaching it to discriminate elements within a logical space of particulars, universals, facts, etc., of which it is already indiscriminately aware, and to associate these discriminated elements with verbal symbols.

2.1. The regress of interpretations

Doubts about Kripke’s account being a proper analysis of Wittgenstein’s views have often and correctly been raised, and yet it established a standard way to carve a skeptical argument in the *Philosophical Investigations*. I would push for a slight but insightful and consequence-laden adjustment of this carving, which I borrow from McDowell (1984).

McDowell pictures the bundle of problems Wittgenstein was dealing with as a complex dilemma: on the one horn, the familiar correspondentistic representation of truth as congruence between meanings and facts, that Wittgenstein rejects, and, on the other horn, the whole famous paradox of §201: «no course of action can be determined by a rule, because every course of action can be made out to accord with the rule». In this adjusted picture Wittgenstein’s argument does not rest in the formulation of a paradoxical

¹ This is slightly adapted from Hattiangadi (2007, p. 56).

conclusion. Rather, it draws a path moving from the rejection of the augustinian correspondentistic interpretation of meanings (the first horn of the dilemma), through the analysis of the difficulties of the paradox of rule-following (the second horn), to their eventual solution. Remarkably, the solution comes just in §201 of *Philosophical Investigations*, when Wittgenstein points out the misunderstanding which supports the whole dilemma: the way the paradox is generated – i.e., by providing explicit definitions for the meaning of a linguistic expression as interpretations of its public usage – preempts the realization that «there is a way to grasp a rule which is *not* an *interpretation*, but which is exhibited in what we call “obeying the rule” and “going against it” in actual cases». Once this misunderstanding is dropped, one can avoid the paradox without committing to any mythical peg where to hang a permanently adequate description of meaning. And this means that skepticism can be defused right from the start.

This is enough for my purposes, so I'll drop McDowell's account here, since I'm not willing to follow him in his analysis of the noninterpretative way in which we grasp meanings. But I need to point out one more crucial thing. As it is well known, Wittgenstein's own conception of the noninterpretative grasp on meanings is a “bedrock” of linguistic behavioral practices. Such a bedrock is “fundamental” to any interpretation of linguistic content, in the sense that our inquiries into the adequacy of our definitions can't dig under it. Let's now try to add this idea to the picture. So, on the one side we reject the idea of self-authenticating non-verbal *data* which could grant the adequacy of a definition of meaning, on the other side we drop the requirement of *justification* for the meaningfulness (the following of a rule) of all verbal episodes: hence we are left with a bedrock of verbal episodes whose meaningfulness doesn't require an interpretation. But here's a worry: granted that our linguistic practices are meaningful and granted that we entertain contents, what can we say about them? Since the bedrock can't be interpreted, how can we represent contents? It seems that we can't do semantics. I turn now to defuse this worry.

2.2. The givenness of the logical space

Very roughly, in Sellars's account, the Myth of the Given is the idea of a prelinguistic sort of epistemic awareness of conceptual contents. This is often construed as the awareness of certain *qualia* – e.g., the preconceptual awareness of *green*, as opposed to *red*. Here, I'm pointing at a more

fundamental level of the Myth: the idea of a prelinguistic awareness of the *determination* – whether it be qualitative or not – of conceptual content. My purpose will be to ensure that, however fundamental the criticism against the Myth may be, it is compatible with a realist stance. In this section, I'll try to show how Quine, the great foe of empiricist dogmas, fell prey of the Myth of the Given. Recall that I need this piece of argument in order to disentangle the thesis that meaning is “inscrutable”, so let me focus on Quine (1968). The outcome of Quine's argument for the inscrutability of reference could be hastily sketched as follows: since “meaning” is too a vague notion for the naturalist, she tries getting a firmer grip on conceptual content by recurring to *extensionalism*, but she has to realize that reference is indeterminate as well. This story is well known, but let me rehearse it very briefly. Consider an extensional semantics in which a domain of objects D is *given* and concepts ϕ , ψ are interpreted as functions whose ranges are collections $K = \{a_1, \dots, a_n\}$ of objects in D , i.e.,

$$x \in K =_{\text{def}} x = a_1 \vee x = a_2 \vee \dots \vee x = a_n$$

and predication is defined as

$$\phi x =_{\text{def}} x \in K.$$

Now, in this framework the analysis of meanings reduces to the inspection of collections K s, whose set theoretical relations generally can be made explicit in terms of quantified sentences. So we *say*, for instance, that rabbits are mammals in terms of

$$\forall x(\phi x \rightarrow \psi x).$$

Going backwards, this is how we inspect the ontology of a theory through the quantificational apparatus. Such is the power of the extensional representation. But, according to Quine, this approach to the analysis of meaning is flawed, since it entirely depends on the givenness of the domain D . For suppose another domain D' was given: collections K' s could be rearranged on D' so to correspond to our concepts and satisfy each and every quantified sentence we used to express their relations. Again, the point is that, although we are provided with all the information expressed in the quantified sentences of our theory, still we can't tell whether ϕx is defined in terms of K or K' . Hence, Quine claims, ontology is *relative* in the sense that it can be specified only with

relation to a given background. The problem is obviously that no background domain is given to us in semantic analysis.²

Up to here I just described a familiar piece of the toolbox of the philosopher of language, now I want to show that it is prone to malfunctioning if it is not properly used. Quine construes the idea that the domain of a theory can be specified only against the background of another given theory in the following way:

We may picture the vocabulary of a theory as comprising logical signs such as quantifiers and the signs for the truth functions and identity, and in addition descriptive or nonlogical signs, which, typically, are singular terms, or names, and general terms, or predicates. Suppose next that in the statements which comprise the theory, that is, are true according to the theory, we abstract from the meanings of the nonlogical vocabulary and from the range of the variables. We are left with the logical form of the theory, or, I shall say, the *theory form*. Now we may interpret this theory form anew by picking a new universe for its variables of quantification to range over, and assigning objects from this universe to the names, and choosing subsets of this universe as extensions of the one-place predicates, and so on. (Quine, 1968, p. 204)

Clearly his concern is that the bare *theory form* provides no information to discriminate suitable models, in the absence of a given background domain where objects could be picked. Now, from a sellarsian perspective, the idea of a «theory form» sounds very suspicious: it amounts to the view that a logical backbone – a formal vocabulary of quantifiers, predicative letters and variables – might be sterilely transplanted from one theoretical body to another. Brandom described this sort of suspect in some detail, by providing reasons against Quine’s famous “gavagai” example.³ Assuming that there is an incompatibility between the sortals the theorist may use to translate “gavagai”, e.g., “rabbit” and “undetached rabbit part”, he must have at his disposal the linguistic resources to make such incompatibility explicit by contrasting two

² It may be (and it has been) objected that no background domain is given to us with the exception of the only one that really guarantees an objective extensional representation of contents: the domain of *all possible* particulars. From the perspective adopted in the present paper, this modal way out might be resisted by noticing that possible world semantics simply does not provide a definition of meanings which is adequate to their usage in linguistic practices, because actual speakers just have no suitable access to such a metaphysical domain. I’ll stick to this rejoinder here for two reasons. First, for the sake of my argument, I don’t want to just cut the gordian knot of adequacy. Second, in an important sense, as I am going to argue, there are *not enough* possible particulars to do the job.

³ See Brandom (1994, pp. 409–412).

general ways to reindividuate items of the different sorts. That is to say, if “gavagai” is a genuine sortal then it must apply either to individual rabbits or to individual undetached rabbit parts, and it must be so applicable in language. So, for instance, it must be possible to say that «“gavagai” *a* is the same as “gavagai” *b*» and check whether *a* can be substituted with *b* in contexts like «... is a mammal» or «... is a broken foot». Surely the theorist has the linguistic resources to do that. The question is: do the natives have similar resources? If they do, then obviously the translation wouldn’t be indeterminate. But if they don’t, then how could a sortal which doesn’t match the natives’ resources for reindividuation be a proper translation?

Notice these reasons do not amount to a refutation yet. Brandom looks at Quine’s example from his already refined inferentialist perspective, and his purpose is to defend *that* perspective from indeterminacy. Quine, from his own point of view, could have easily dismissed the question about the use of sortals by recurring to extensionalism. In fact, what guarantees the feasibility, *in principle*, of the procedure described by Quine in the above quotation is just the extensional definition of concepts: in order to see whether ϕ is incompatible with ψ it is enough to check whether K and K' are disjoint in DUD’.

But then, is there something wrong with extensionalism? This point is a tricky one. Sellars explicitly tackled it at the very end of Sellars (1957, §§ 102–108). Now, that is a mouthful of sellarsian philosophy, but here I’ll try to chew just the very bit I really need:

[T]he logic of variables and quantification involves not only the *momentary* crystallized content of the language at a cross section of its history, but also its character as admitting – indeed demanding – modification, revision, in short, development, in accordance with rational procedures. (Sellars, 1957, § 105)

Sellars makes two crucial claims here: (a) the extensional description of contents pictures an idealized phase of linguistic practices in which conceptual resources are completely and definitively made explicit; (b) such a description implies, indeed requires, the possibility of content to be continuously improved. I won’t be able to say anything about (b) until the end of this paper, so, for the time being, I’ll focus on (a).

This remark of Sellars’s highlights what goes wrong with Quine’s strategy: one just can’t recur to the quantificational apparatus as a sterile scalpel to carve ontologies out of theoretical bodies and compare them, because the

quantificational apparatus itself is part of those bodies. What does that mean out of metaphor? Consider a standard tarskian semantics and try to ask, for instance: what is a model for a sentence p ? It seems that we can provide two answers: (i) a model for p is an *interpretation* that gives to p a designated semantic value; (ii) a model for p is a *possible world* that contains the state of affairs represented by p . Indeed, the fact that (i) and (ii) can often be treated as equivalent is the key of the success of model-theoretic semantics. But that doesn't mean we can simply equate them: on the one side we have the interpretation of linguistic expressions, on the other side we have the representation of states of affairs. My purpose here is not to highlight a gap between (i) and (ii). To the contrary, I want to claim that since there is no gap and we make explicit the interpretation of language in terms of an extensional representation, we can't account for the variability in the determination of contents inside this representation: we'll never get different interpretations of the same concept, we'll always get just different concepts. In this sense, Quine fell prey of the Myth of the Given to the extent that he took the quantificational apparatus of a language as conceptually prior to the determination of its contents. However, and this is what Sellars's remark (b) is about, it is just because we provide an explicit interpretation of contents that we can modify them. So we can't just throw away semantic theory.

3. A theory of meaning

In the previous section I tried to open up some space for the possibility to endorse both a realist stance in semantics and a non-reductive account of the normativity of meaning. Now I want to substantiate this possibility. Let's start from scratch once again by asking: what do we do, when we do formal semantics?

3.1. Picturing and meaning

First of all it must be acknowledged that the formal semantics we've all learned, model-theoretic semantics, is couched in the representational tradition. Intuitively, the paradigm of representation is a *picture*. But consider a prototypical statement like "the particle a has spin- s ", and ask: what does it *represent*? The answer, clearly, will be that it represents the particle a as having a certain spin. Now, in a sense this means that the statements *says of a that it*

has spin-*s*. In another sense it means that the statement *pictures* a *complex object*, e.g., the *fact that a* has spin-*s*, in the same way as a plan represents a building (where a building is a complex object and a single brick is not). This latter sense is quite misleadingly supported by the formalization of the statement in first order logic, say “*S(a)*”: $*S(a)*$, as a sign design, is a complex object as well, so that $*S(*)*$ and $*a*$, as sign designs, are *part of* it. In this latter sense our statement is construed as a complex name for the complex object it pictures.

Sellars repeatedly denounced the risk that the confusion of these two senses may engender the confusion between asserting and picturing. In particular it leads to think that the representational purport of linguistic expressions can be explained in terms of a picturing relation holding between them and things they *name* in the world: indeed, while linguistic expressions do, in a sense, picture things in the world, it is this very fact that begs the question about how *representings* can point beyond themselves to *representeds*.

3.2. Linguistic roles

So, we must also acknowledge that when we do formal semantics we do not, in the first instance, *describe* a relation between a language and the world. But then, what are meaning statements about? Consider again:

“ ϕ ” (in L) *means* F

Let me lay down some platitudes. First, “ ϕ ” on the left, a quotation of the sign design $*\phi*$, is the name of an expression in language L. Second, F on the right can’t be the name of a linguistic expression in L as well, on pain of regress. Third, however, if F were not a name then either it would not occur in a truthfunctional sense or *meaning* would not be a relation (but then what would it be?). So F must be the name of something, and it must be something we have some knowledge of in order for meaning statements to be explanatory. Indeed, F is usually construed as an expression of *our* meta-language. But what does it mean to have some knowledge of an expression of our language? It means that we know how to use it: we know how to deploy it with relation to the context, we know its *linguistic role*. Linguistic roles provide a functional classification

of linguistic sign designs. If we adopt Sellars's dot-quotation⁴ to single out these roles, we can provide an alternative interpretation of meaning statements:

“ ϕ ” (in L) means $F \rightarrow$ “ ϕ ” (in L) is an $\bullet F \bullet$.

Distributive terms like $\bullet F \bullet$ designate (both intra- and inter-linguistically) expressions which are used in the same way, that is to say, those expressions which occupy the same place in the net of relations – paradigmatically, inferential relations – established on linguistic sign designs by their use in linguistic practices. Notice that it's easy to accommodate in this framework the semantic realist's interpretation of linguistic rules, since part of the linguistic role of $\bullet F \bullet$ s might be, for instance, to apply whenever $\bullet f \bullet$ s apply as well. Notice also that all this is compatible with representational semantics – indeed, this is purported to be just its correct interpretation.

So what we do when we do formal semantics is to provide a model for linguistic roles in order to *explain* the use of linguistic sign designs. Where the world comes into the picture is in the evaluation of the model. As it was clearly stated just in Tarski (1944), formal semantics is quite independent from the ideas we may entertain about the *nature* of semantic contents, but, obviously, that doesn't mean that formal semantics doesn't explain anything about language and mind. In fact it's easy to see how models work here. In Section 3.1 we've noticed that we can talk of sentences as complex sign designs that *picture what it is*, without implying that this fact should explain how they *say of what it is that it is*. We've acknowledged that the explanation runs in the opposite direction: it is because we use certain sign designs according to certain rules in order to say something, that those sign designs picture it (rather than something else). So, semantic models do not apply to a picturing relation between expressions and things in the world, because that would not

⁴ Just to briefly sum up, dot-quotation applies to expressions in a given *familiar* language to build *distributive singular terms* referring to any expression in any language that play the same linguistic role of the quoted expression. So, as the *distributive singular term*

the pawn

refers to any *piece* (however materially realized) that is subject to certain rules in a chess game, in the same way the *distributive singular term*

\bullet triangular \bullet

refers to any sign design (however linguistically realized) that is subject to certain rules in a language game. For further details see Sellars (1963a).

provide an explanation of linguistic sign designs. That is, to repeat, meaning statements do not establish relations between linguistic items and non-linguistic items. Instead, semantic theory provides models which represent *linguistic roles* in order to explain the use of linguistic sign designs.

Thus, to take the realist stance as to shortcircuit the explanation of the contents by interconnecting the linguistic domain of sign designs with the non-linguistic domain of the model is just a mistake.

3.3. The realist scruple

The notion of linguistic role might seriously worry the naturalistic biases of the semantic realist, and she may be willing to protest that I'm just weighing her down with ontological burdens while refusing to answer the only relevant semantic question: how are linguistic expressions related to things in the world?

What I would be missing can be best highlighted by appealing to common sense. So, suppose one morning I wake up with a terrible rash. I go to the doctor who, after checking me up, declares: "you have chicken pox!" So, I adopt all the necessary cures, which involve, according to Italian lore, to devote myself to Saint Anthony. My behavior ensues from the fact that "chicken pox", in the doctor's claim, means *chicken pox*, rather than, for instance, *encephalitis*, and "you" means *me*. Otherwise I would have had to devote myself to Saint Paul, or simply do nothing at all.⁵ In other words, the complaint is that I'm wavering on the notions of *Reference* and *Truth*. But, as a matter of fact, it is just at this point of the analysis that such notions can be properly introduced. And the reason is that we can now see clearly what they are not. On the other hand, a proper account, as it's easy to realize, would greatly exceed the space of this paper, so I'll have to be very schematic.

Let's begin with *Truth*: it is not a relation between sentences and states of affairs, nor a relation among sets of propositions. The isomorphism between the structure of linguistic sign designs and the structure of states of affairs is a necessary yet not sufficient requirement for the picturing relation between language and world to hold. If the meaning of a propositional sign design is its propositional linguistic role, then to say that a sentence is true is to say that it is *correctly* used according to its role – paradigmatically, when it is correctly

⁵ The example is ironically adapted from the introduction of Hattiangadi (2007).

asserted according to the inferential rules of the language. Notice that these inferential rules involve the *context*, and, in particular, the consideration of both possibly perceptual premises and possibly practical consequences of the assertion: it is crucial, in order to understand how the uses of expressions get in touch with the world, to take into account in this sense full-fledged linguistic practices. Notice also that these rules establish the validity of not only *formal*, but also *material* inferences. The sellarsian distinction between formal and material inferences is defended in Sellars (1953). Here a concise but harshly bolzanian definition could be put as follows: formal inferences are valid because they are substitutional occurrences of patterns in which “logical” vocabulary is fixed, while material inferences are valid because they are substitutional occurrences of patterns in which “non-logical” vocabulary is fixed.

Now, all this holds for subsentential expressions as well, so, in a sense, this is also account of *Reference*. But such an account, as it stands, would be incomplete in two main aspects. First, subsentential roles are not all of the same sort. So, just to get the idea, consider again the sign design $*S(a)*$: part of “*a*” being an $\bullet a \bullet$, as contrasted to being an $\bullet S(\) \bullet$, is that “*a*” is a singular term. This difference must be accounted for in some way, and it must be explained how the linguistic role for sentential sign designs like $*S(a)*$ is composed by the linguistic roles of $\bullet S(\) \bullet$ s and $\bullet a \bullet$ s.⁶ Second, our use of sortals elicits regularities which pertain to our concepts of thing-kinds (or “essences”). The way in which the linguistic role of sortals establish those regularities must be accounted for as well. In Sellars (1957) that is done in terms of the modally robust, and yet defeasible, inferential rules which define the linguistic role of sortals. So, for instance, the inferential rules which define the linguistic role of $\bullet \text{match} \bullet$ may involve conditionals like

⁶ This point has been often construed as revealing the main technical flaw of the whole picture: the semantic structure and the syntactic structure have to be isomorphic, such is the requirement of *compositionality*, but linguistic roles do not compose. However, this is wrong. To begin with, Chapter 6 of Brandom (1994) shows how to evaluate the inferential role of subsentential expressions by exploiting just compositionality. Still, some doubts are raised by the exploitation of compositionality in an holistic framework (Fodor and Lepore, 2001, 2007): the proof of inferential conservativity for the logical vocabulary of *Incompatibility Semantics* in Brandom (2008) shows how they can be dismissed. In this respect, the general point to be realized is that, from a formal point of view, compositionality is but a trivial problem, as it is clearly explained in Westerståhl (1998). More interesting issues rise when the semantic and syntactic structures are already defined according to other theories and independent assumptions.

“This match would light if scratched”,

i.e., they entitle one to move, defeasibly, from the application of a **•match•** and a **•scratched•** to an application of a **•lighted•**.

To sum up, semantic notions like *Reference* and *Truth* are *normative* notions, in the sense that they are employed to specify correct use.

Here, a common misunderstanding should be dismissed straight away. As a general objection against this analysis of meaning in terms of rules, one may notice that it is just a mistake to try to define semantic notions in terms of *prescriptions* on the use of expressions, since these latter have to do with the pragmatics of language.⁷ This worry is legitimate, but, *qua* objection, it loses all its bite as soon as it is noticed that it is grounded on the confusion that consists in interpreting rules *directly* as patterns of behavior. Behavior involves *actions*, and rules for linguistic actions are obviously of the pragmatic sort, like: *tell the truth!* Now, while linguistic episodes can (and usually do) manifest themselves in terms of overt verbal performances, the rules that define their contents in the sense here intended are not, to use a Sellarsian turn of phrase, rules of performance, but rules of *criticism*. That is, they do not specify what one *ought to do*, but what one *ought not*. So are, for instance, the rules of the Highway Code: they do not say where one ought to go, but they say that one ought not to cross a red light. Since linguistic rules are usually construed as inferential rules, I’m afraid that sometimes this confusion may be backed by a certain hasty suggestion about the notion of *following logically*, a suggestion already mocked by Lewis Carroll but rather die hard, according to which a logical inference, e.g., from *p* to *q*, amounts to a *prescription* of asserting *q* once *p* is asserted (or, at least, to believe *q* once *p* is believed): this, as the tortoise tried to explain to Achilles, is obviously nonsense.

4. Normative vocabulary

Almost every piece of my argument is in its place, and still, it seems, my results are quite poor. Even if I was successful in the criticism of the standard picture of the debate about normativity of meaning, I haven’t yet indicated any other practicable way to get off the ground in semantics. In order to provide an alternative I need to introduce, as the last bit of my argument, the analysis of the role of normative vocabulary.

⁷ See for instance Wikforss (2001), Hattiangadi (2007), Glüer and Wikforss (2009).

The fundamental role of normative vocabulary eventually comes to the fore in the attempt to cash out the notion of *correctness* established by the rules which define linguistic roles. In fact, if we were to ask how to evaluate the correctness of a particular application of the rules, immediately we would have to face the same old regress arising again and again. However, the key to defuse the regress once and for all lies just in a proper analysis of normativity.

Before we go ahead, let me briefly recapitulate. Formal semantics provides models to explain linguistic behavior, but this theoretical enterprise seems to be bound to failure because of the irreducibility of the normative character of human behavior to descriptions of matters of fact – such is the wittgensteinian image of the bedrock. In this paper we have established (i) that neither the idea of such a bedrock nor the possibility of its theory is incoherent, and (ii) that, once a functional characterization of contents is accepted, the realist stance is compatible with a theory of linguistic roles. Still, it must be clarified how the normative analysis of linguistic roles may fit into the explanation of linguistic behavior provided by formal semantics. In other words, through the functional characterization of conceptual contents in semantics we may obtain scientific theories whose models explain our linguistic behavior, but still fall short of explaining it as rational behavior.

But now I have all I need to start the last bit of my sellarsian analysis. As it's easy to notice, the problem we face here is but a particular occurrence of one of the major themes of Sellars's: he described it in Sellars (1962) as the problem of fusing in a "stereoscopic vision" the "manifest" and the "scientific image of man in the world".⁸ I won't try to approach a proper analysis of the general theme, but I hope that my account of this particular instance was clear enough to let the reader understand how Sellars's solution applies here. Thus, the key to obtain such stereoscopy is the following:

[T]o complete the scientific image we need to enrich it *not* with more ways of saying what is the case, but with the language of community and individual intentions. (Sellars, 1962, p. 78)

Here the notion of community intention is crucial, so it's worth taking a break and making it clear. First, the notion of intention is, as expected, a normative notion. In this context, intentions are not dispositions to act. Rather, they are the sort of things which can move rational beings, those who are sensitive to the force of reasons. Thus, for instance, it is *because* you have an intention

⁸ See Sellars (1962, pp. 40–41).

whose content is expressed by “I shall raise my arm” that you raise your arm. As such, intentions are evaluated in terms of *reasonableness* with relation to practical reasoning. For instance, we express the fact that my intention (not) to raise my arm is reasonable in circumstances of kind *C*, by saying “I ought (not) to raise my arm, in circumstances of kind *C*”. Second, the evaluation of the reasonableness of intentions poses the problem of *objectivity*. Notice that this might be a harmful trigger for the regress of interpretations. However, and this is the crucial point, at the level of intentions, a direct solution to the regress can be provided.

Let’s begin by asking what it means to generalize subjective intentions. The idea is to move from principles of the form

I ought to do actions of kind *A*, in circumstances of kind *C*

to principles of the form

Anyone ought to do actions of kind *A*, in circumstances of kind *C*.

Such a generality can’t be achieved by intentions of the form

Everyone shall do actions of kind *A*, in circumstances of kind *C*,

for we know that moral principles, *qua* normative principles, are not reducible to what everyone does. So our question turns out to deal with the analysis of

Anyone shall do actions of kind *A*, in circumstances of kind *C*.

Here, Sellars suggest, “anyone” refers to any of those who *share* the intention. To share an intention in this sense is not to have the *same* subjective intention, but to have a community intention of the form

We shall do actions of kind *A*, in circumstances of kind *C*.

It is the sharing of intentions that constitutes a community.⁹ The constitution of a community of rational agents who share intentions supports the application of normative values to every domain pertaining to their agency, in particular to the epistemic and the practical domains. It is in this sense that Sellars’s solution to the problem of fusing the images consists in *extending* the descriptive vocabulary of theoretical representations of the world by introducing the normative vocabulary of communities of rational agents.

Now let’s go back to our argument. In order to accept Sellars’s suggestion here we have to block two preliminary objections.

⁹ See Chapter VII of Sellars (1968).

The first objection is a rehearsal of a worry we already considered in Section 3.3: does this approach commit us to intentions *qua* mental sort of things? The answer is no, because we only need to accept into our models theoretical objects which can be defined as functional analogues of sign designs. This, notice, entitles us to dismiss the hardest part of Sellars's own problem, i.e., the reception of *persons* inside the scientific image: since we are interested in formal semantics, we have just to deal with the functional analysis.

The second objection may be formulated as a request for clarification: Kripke already invoked the perspective of the community but failed to establish it as anything more than an intersubjective point of view, so how is Sellars's proposal different? The answer is that the two proposals are indeed pretty different: according to Kripke, the community establishes the horizon of the generality of norms, while, according to Sellars, the universality of norms establishes the horizon of the community. In this sense, the very task of making the contents of the norms explicit is the task of constituting the community of the agents who follow them. This is probably the hardest point to acknowledge, but it's also probably the most important: surely, without it the whole approach would be idle.

Then, the application of Sellars's solution to the semantic domain is quite straightforward. In fact, *Truth* turns out to be the sort of reasonableness of community intentions which pertain to the semantic domain, and its evaluation has to be made explicit in terms of normative principles pertaining to linguistic performances. Typically these will be principles of the form

One ought not to refuse to apply " ψ ", if one accepts to apply " ϕ ".

These principles establish relation among the sign designs – e.g., $*\phi*$, $*\psi*$ – which are employed in the practice they regulate. In this sense each sign design is identified by the *functional role* – e.g., $\bullet\phi\bullet$, $\bullet\psi\bullet$ – it acquires in these relations. The goal of formal semantics is to represent these linguistic roles in terms of relations among the elements of the model of a semantic theory.

5. Concluding remarks

I'm afraid that I haven't put forward any really new thesis in this paper. Or at least, it seem to me that I've just collected pieces of reasoning which had been already in good sight there on the table, although they are often ignored in the debate about the normativity of meaning. My attempt here was to tidy things up a bit: sometimes that is enough to put them back into good use. Thus, I hope

that the alternative realist stance was stated neatly enough. In this last section I want to suggest some reasons to adopt it. I already pointed out its assets with relation to some traditional problems in philosophy of language, so what I still have to show is that it is suitable to satisfy the prior needs of the realist. In other words, the question I still want to ask is: is it really a *realist* stance? Let me pick a paradigmatic list of the *desiderata* for a realist stance just from a scholar who aims to defend semantic realism against skeptical attacks. In the introduction of Hattiangadi (2007) the position of the semantic realist is basically summarized in two tenets: (a) to understand the meaning of an expression is to grasp its correctness conditions; (b) ascriptions of meaning are subject to correctness conditions as well.

Indeed, the approach presented in this paper satisfies both these *desiderata*. With respect to (a), the present approach takes meanings to be defined by the inferential rules of a language and to be represented in terms of linguistic roles: according to the present approach to understand the meaning of an expression is to grasp its linguistic role, which represents the conditions for its correct application. Therefore the present approach satisfies tenet (a). With respect to (b), the present approach takes the objectivity of the rules of language to be defined by the horizon of the community of speakers who share the same normative space: according to the present approach the correctness of meaning ascriptions is evaluated in terms of the linguistic norms of the community of speakers. Therefore the present approach satisfies tenet (b).

In spite of this, the semantic realist might still feel unsatisfied. The reason is that there is a third idea implicit underneath tenets (a) and (b), the idea that correctness conditions have to be checked against some sort of “fact of the matter”, as Hattiangadi suggests. Depending on how the notion of “fact of the matter” is construed, this idea may turn out to be a third tenet (c) characterizing the semantic realist’s position. There is a sense in which the present approach satisfies tenet (c) as well. The present approach takes formal semantics to picture states of affairs by representing semantic rules pertaining to them, so that according to the present approach there’s a precise sense, formally specifiable in terms of semantic models, in which correctness conditions are checked against facts of the matter. In this sense, however, tenet (c) is already contained in (a) and (b). There is obviously another way to maintain (c). It consists in construing the notion of “fact of the matter” as a given source of conceptually determinate information available out there (or in here) to be grasped. It’s just one major credit of Sellars’s that of having

delivered to contemporary philosophers the consciousness that such a given source is a myth. In this latter sense (c) is an autonomous tenet, which, however, I think, doesn't (and shouldn't) belong to a characterization of semantic realism. Actually, in this latter sense (c) has often been construed as marking the boundary between realism and antirealism. In fact, it's always been the antirealist strategy, at least since Gorgias, that of denying (a) and (b) by casting doubts on the existence of objective facts of the matter. But, as we have seen, it is just a mistake to believe that objectivity is given independently of normative practices, and the semantic realist has no need and no gain in following her opponent in this mistake.

I wish to conclude by adding that the sort of semantic enterprise envisaged at the end of the previous section is not is not just wishful thinking. On the one side, the inferential analysis of meaning is an ongoing logical enterprise being presently developed, among others, by Jaroslav Peregrin and by Dag Prawitz and his followers.¹⁰ On the other side, the account of normativity providing the interpretation for inferential relations has reached some practicable results, mainly due to Robert Brandom's elaboration of sellarsian themes.¹¹

REFERENCES

- Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- Brandom, R. (2008). *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford: Oxford University Press.
- Fodor, J., & Lepore, E. (2001). Brandom's Burdens: Compositionality and Inferentialism. *Philosophy and Phenomenological Research*, 63(2), 465–481.
- Fodor, J., & Lepore, E. (2007). Brandom Beleaguered. *Philosophy and Phenomenological Research*, 74(3), 677–691. Reprinted in Weiss & Wanderer (2010).
- Glüer, K., & Wikforss, Å. (2009). Against Content Normativity. *Mind*, 118(469), 31–70.

¹⁰ See Prawitz (2006) and Peregrin (2009, 2010).

¹¹ See, again, mainly Brandom (1994, 2008).

- Goodman, N. (1954). *Fact, Fiction & Forecast*. London: University of London.
- Hattiangadi, A. (2007). *Oughts and Thoughts: Rule-Following and the Normativity of Content*. Oxford: Oxford University Press.
- Kripke, S. (1982). *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.
- McDowell, J. (1984). Wittgenstein on Following a Rule. *Synthese*, 58, 325–364. Reprinted in McDowell (1998).
- McDowell, J. (1998). *Mind, Value, and Reality*. Cambridge, MA: Harvard University Press.
- Peregrin, J. (2009). Inferentialism and the Compositionality of Meaning. *International Review of Pragmatics*, 1(1), 54–181.
- Peregrin, J. (2010). Inferentializing Semantics. *Journal of Philosophical Logic*, 39(3), 255–274.
- Prawitz, D. (2006). Meaning Approached Via Proofs. *Synthese*, 148(3), 507–524.
- Quine, W.V.O. (1968). Ontological Relativity. *Journal of Philosophy*, 65(7), 185–212. Reprinted in Quine (1969).
- Quine, W.V.O. (1969). *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Sellars, W. (1953). Inference and Meaning. *Mind*, 62(247), 313–338. Reprinted in Sellars (1980) and Sellars (2007).
- Sellars, W. (1956). Empiricism and the Philosophy of Mind. In H. Feigl & M. Scriven (Eds.), *Minnesota Studies in the Philosophy of Science*, vol. I. Minneapolis: University of Minnesota Press, 253–329. Reprinted with additional footnotes in W. Sellars (1963), *Science, Perception and Reality*. New York: The Humanities Press, 127–196. Republished as a volume (1997) with an introduction by R. Rorty and a study guide by R.B. Brandom, Cambridge, MA.: Harvard University Press.
- Sellars, W. (1957). Counterfactuals, Dispositions, and the Causal Modalities. *Minnesota Studies in The Philosophy of Science*, II, 225–230. Reprinted in part in E. Sosa (Ed.) (1975), *Causation and Conditionals*.

Oxford University Press.

- Sellars, W. (1962). Philosophy and the Scientific Image of Man. In R.G. Colodny (Ed.), *Frontiers of Science and Philosophy*. Pittsburgh: University of Pittsburgh Press, 35–78. Reprinted in Sellars (1963b) and in Sellars (2007).
- Sellars, W. (1963a). Abstract Entities. *Review of Metaphysics*, 16, 627–671, Reprinted in Sellars (1967); reprinted in Sellars (2007).
- Sellars, W. (1963b). *Science, Perception and Reality*. New York: The Humanities Press. Reprinted by Ridgeview Publishing Company, Atascadero, California, 1991.
- Sellars, W. (1967). *Philosophical Perspectives*. Springfield, ILL: Charles Thomas. Reprinted as : W. Sellars (1977). *Philosophical Perspectives*, 2 voll. (vol.1: *History of Philosophy*; vol. 2: *Metaphysics and Epistemology*). Atascadero, CA: Ridgeview Publishing Co.
- Sellars, W. (1968). *Science and Metaphysics: Variations on Kantian Themes*. New York: Humanities Press.
- Sellars, W. (1980). On Reasoning About Values. *American Philosophical Quarterly*, 17(2), 81–101.
- Sellars, W. (2007). *In the Space of Reasons: Selected Essays of Wilfrid Sellars*. Cambridge, MA: Harvard University Press.
- Tarski, A. (1944). The Semantic Conception of Truth and the Foundations of Semantics. *Philosophy and Phenomenological Research*, 4(3), 341–376.
- Weiss, B., & Wanderer, J. (Eds.) (2010). *Reading Brandom: Making it Explicit*. London: Routledge.
- Westerståhl, D. (1998). On Mathematical Proofs of the Vacuity of Compositionality. *Linguistics and Philosophy*, 21(6), 635–643.
- Wikforss, Å. (2001). Semantic Normativity. *Philosophical Studies*, 102(2), 203–226.
- Wittgenstein, L. (2009). *Philosophical Investigations*, 4th edition. Trans. and ed. by P.M.S. Hacker and J. Schulte. Oxford: Blackwell.