

Why compositionality?

Jeroen Groenendijk & Martin Stokhof*

G.N. Carlson & F.J. Pelletier (eds), *Reference and Quantification. The Partee Effect*.
CSLI Publications, Stanford, 2005, pp. 83–106

Abstract

The paper identifies some background assumptions of compositionality in formal semantics and investigates how they shape formal semantics as a scientific discipline.

1 Introduction

In the enterprise of formal semantics, as it grew out of the pioneering work of Montague, Davidson, Lewis, Hintikka and others, compositionality occupies a special place. On the one hand it has been identified, correctly in our opinion, as one of the most central principles of formal semantics. However, far from being taken for granted, it has been, and still is, discussed frequently from a variety of perspectives. This is an indication, we venture, that all is not well in this sense that semanticists intuitively are aware that the role that the principle of compositionality plays in their discipline is in need of further clarification.

That compositionality occupies a central position is not remarkable as such. It is a principle that concerns the relation between syntax and semantics, and since there can be no semantics without syntax, it is at the very heart of the enterprise. Informally, compositionality holds that the meaning of a complex expression is determined by the meanings of its parts, and this, admittedly vague, characterisation makes clear that both syntax (structure, i.e., expressions and their part–whole relationship) and semantics (meaning, i.e., the semantic values of expressions and their part–whole relationship) are involved. This also makes clear that an assessment of the status of the compositionality principle is intimately tied to whatever view we have of both structure and meaning.

This interdependence between the status we accord to compositionality and our views on syntax and semantics is acknowledged in what we may characterise as the ‘received view’ concerning the principle among formal semanticists. It considers compositionality as having a theoretical status. Unless we have independently motivated constraints on syntax and/or semantics, i.e., particular

*ILLC / Department of Philosophy, Universiteit van Amsterdam. We would like to thank an anonymous referee for helpful comments on an earlier version.

do's or don't's concerning the nature of syntactic structure, constituenthood, rule formats, or concerning the nature of meanings, their internal structure and possible interactions, the principle of compositionality is not an empirical hypothesis. Rather, it must be viewed as a methodological principle, one that represents a choice to do semantics in a particular way. The following quotation from Janssen may serve as an illustration:

The most valuable arguments [for compositionality, G&S] are, in my opinion, those concerning the elegance and power of the framework, its heuristic value and the lack of a mathematically well defined alternative. [...] I expect that problematic cases can always be dealt with by means of another organisation of the syntax, resulting in more abstract parts, or by means of a more abstract conception of meaning. The principle only has to be abandoned if it leads too often to unnecessarily complicated treatments.¹

Its justification, hence, is not empirical, but rather resides in the 'fringe' benefits of doing semantics in this way: elegance, perspicuity, a precise mathematical formulation. Of course, it is not to be ruled out that in some cases the compositional alternative is less intuitive than its non-compositional counterpart.

According to the received view, it is only when certain constraints are operative, for example concerning the format of the syntactic rules one may employ, or concerning the kinds of mathematical objects that one allows playing the role of meanings, that the question of compositionality becomes an empirical one. As Barbara Partee formulated it:

... the principle can be made precise only in conjunction with an explicit theory of meaning and of syntax, together with a fuller specification of what is required by the relation 'is a function of'. If the syntax is sufficiently unconstrained and meanings are sufficiently rich, there seems no doubt that natural languages can be described compositionally. Challenges to the principle generally involve either explicit or implicit arguments to the effect that it conflicts with other well-motivated constraints ...²

When additional, independently motivated constraints on syntax and/or semantics are formulated, compositionality can be treated as a property that syntactic constructions and the corresponding meaning assignments may or may not have. Then the question whether the complex expressions that result from a particular mode of syntactic construction can be assigned meanings in such a way that compositionality is satisfied, is viewed as a meaningful, empirical question.³ However, it should be noted that the purported empirical character of compositionality essentially depends on the nature of these independently motivated constraints: if the latter are not empirical then neither is the question

¹Cf., [12, pp. 38, 39]. Janssen's pioneering thesis contains an extensive defence of this view.

²Cf., [16, p. 281].

³Cf., [16, 13] for extensive discussion of various kinds of examples.

of compositionality. In particular when it comes to the nature of semantics, one seems hard pressed to come up with general principles that are truly empirical in nature. Of course, observation of empirical facts constrain the analysis, and in some cases the tools one uses, but that is another matter. So it seems the methodological view sticks, at least to some extent.

Another role that the principle of compositionality plays and one that seems to be a natural development of the methodological view, is that of an arbiter between rival analyses of a given set of phenomena. Here one uses compositionality as a measure by which analyses are rated, usually with a preference for the (more) compositional ones. This same perspective also steers the development of semantic analyses. A good case in point is the development of various semantic theories as compositional alternatives for discourse representation theory.⁴ As said, this role of compositionality aligns more with the methodological than the empirical view. But the motivation is not (always) simply that ‘since it can be done compositionally, it should be done that way’. For example, when dynamic predicate logic was formulated as a compositional alternative for discourse representation theory, one of the main motivations was philosophical. In dynamic predicate logic one wishes to avoid the assumption of a level of representation in between syntax and meaning that discourse presentation theory needs to make, so as to be able to remain neutral with respect to the question whether any such level of representation is essential in semantics.⁵

As we said, the above seems to represent something of a standard view. That is not to say that it has not been challenged, questioned, modified. Strict and less strict versions of the principle have been proposed and compared; the impact of various explications of the concepts of ‘structure’ and ‘meaning’ on its plausibility have been studied; the existence of apparent exceptions, such as idioms, has been investigated; the technical aspects of its implementation in a grammar model have been scrutinised extensively; attempts have been made to formally prove its vacuousness; the historical antecedents of the principle have been subject of much debate.⁶ Nevertheless, it seems that what we have called the standard view is correct also in this sense that compositionality is a methodological standard by which formal semanticists go about their business: describing and explaining semantic phenomena by devising formal models.

*

What follows is not intended as another contribution to these debates. Rather what we want to do in this paper is take a step back and reflect on the role of compositionality in formal semantics as such. First we ask why it is that compositionality has come to play this role, and we will identify two basic assumptions concerning language and competence that seem to be involved. Next

⁴As was noticed by Szabó, this use of compositionality in fact hinges on the notions of ‘meaning’ and ‘structure’ being polyvalent (cf., [21, p. 479–80]).

⁵Cf., [8, p. 93–8]. The connection between compositionality and representations was already made by Partee (cf., [16, p. 286–7]).

⁶Cf., among many others, [26, 22, 14, 21, 11, 23].

we will investigate whether a change in perspective on exactly what the content of these notions is, might bring about a change in the role of compositionality as well. Finally, we will consider what are the consequences, if any, of these observations for the question what kind of discipline semantics is.

2 Assumptions of compositionality

When we reflect on why compositionality has been such an issue it appears we can discern at least two main reasons: the use of logical languages as representational devices, and the notion of ‘creativity’ of language, or of language users. Let us investigate both in some more detail.

That the use of a logical language in semantics, be it that of first order logic or that of intensional type theory or some other language, brings along the issue of compositionality seems rather obvious. After all, the idea of giving a compositional interpretation of these logical languages goes back to their very origin, in the work of Frege, Peirce, and others. Compositionality is often called ‘Frege’s Principle’, and although there is ample room for debate as to whether the particular way in which it is used today is one that Frege himself would have subscribed to,⁷ there is no doubt that the idea as such is implicit in the way in which logical languages de facto are set up and used. The people who started developing the idea of a logical semantics for natural language, took their lead from the way in which the syntax and the semantics of formal languages go hand in hand and made the principle of compositionality one of the cornerstones of the enterprise. As for the issue of its actual historical origins, that is really an altogether different matter. Correctly or not, they attributed compositionality to Frege, and given the prevailing, rather ‘stylised’ conception of the history of the discipline, one has to admit that this attribution presented itself as obvious.

Within this first main reason why compositionality became a key notion in the semantics of natural language we may discern two elements. First of all, compositionality plays a central role in formal languages as such, relating their semantics to their syntax in a perspicuous way.⁸ Secondly, these very formal languages were considered as models for natural languages and consequently used as auxiliary devices used in formulating the semantics of the latter.⁹ The

⁷Cf., [11], already referred to above, for a recent overview of the state of the debate.

⁸Which is not to say that compositionality always applies rigorously. In some cases it does not and the reformulations that are necessary to show that it can be made to after all, also show why a (moderately) non-compositional alternative is preferred. A case in point is the compositional re-interpretation of quantifiers in first order logic, which requires a semantics with a much more complicated mathematical structure than the familiar one. Cf., [13, p. 422 ff.] for details.

⁹This assumption is quite nicely illustrated by the following quotation from Montague’s seminal paper ‘Universal Grammar’ ([15]):

There is in my opinion no important theoretical difference between natural languages and the artificial languages of logicians; indeed, I consider it possible to comprehend the syntax and semantics of both kinds of languages within a single natural and mathematically precise theory. On this point I differ from a number of philosophers, but agree, I believe, with Chomsky and his associates.

result is that compositionality is transferred quite naturally as a feature, or a requirement, from formal languages to natural languages. So, one reason why compositionality is central to natural language semantics is that we use formal languages as models and representational devices and that the latter are compositional. Which suggests that things might be different if we attribute a different, or perhaps even no role to formal languages.

Also, we might ask ourselves to what extent a formal language can be used as a model for a natural language: Are there not obvious differences, for example in the way they originate, change and develop, in the ways they are used in various circumstances? Here we should note that the Montagovian position so concisely and clearly formulated in the passage from ‘Universal Grammar’ that was quoted in footnote 9 actually is a subtle one: according to Montague there is ‘no important *theoretical* difference’ [our emphasis] between natural and formal languages. This obviously leaves room for ‘non theoretical’ differences. What these might be, we must guess, but we can be sure that according to Montague, whatever they are, they do not make any essential difference for what meaning is and how it can be analysed. That Montague makes clear when he continues that according to him ‘it is possible to comprehend both the syntax and semantics’ of both formal and natural languages by essentially the same means. That implies that any difference between formal and natural languages that is related to the fact that the latter originate and are learnt in a different fashion, are put to many a different use, are constantly changing and adapting, can not make any substantial difference for what meaning is and for the way in which expressions have meaning. Whether that is an acceptable assumption is something we will discuss later on.

Let us now turn to the second assumption that was mentioned above, the one concerning creativity (or ‘productivity’ as it is sometimes called). Sometimes this is considered as a feature of natural languages as such, but it is also an ability that is attributed to language users. However, it seems that these two ways of formulating what creativity is, are regarded as mere variants of each other.¹⁰ This actually illustrates an interesting and important, though seldom noted, assimilation of two sets of notions: those pertaining to language and those related to linguistic competence. As the case of creativity nicely exemplifies, core properties of linguistic competence and distinctive features of language are supposed to mirror each other in such a way as to become exchangeable, almost indistinguishable.¹¹ That such a straightforward identification is problematic is easy to see. For one thing, ‘creativity’ is actually a catch all phrase that covers

¹⁰Cf., e.g., the following well-known passage from Frege’s late paper ‘Compound Thoughts’ ([6]):

It is astonishing what language can do. With a few syllables it can express an incalculable number of thoughts, so that even a thought grasped by a terrestrial being for the very first time can be put into a form of words which will be understood by someone to whom the thought is entirely new.

Obviously, for Frege creativity as a feature of language is mirrored by an ability of its users.

¹¹Another, more contemporary illustration of how the two are run together is provided by the following quote from Davidson’s seminal paper ‘Truth and Meaning’ ([2]):

a number of phenomena that may or may not have much more in common than being called by that name. In modern linguistics it is most often associated with the existence of recursive rules in syntax, and hence mainly relates to the potentially infinite number of syntactic structures and associated meanings that a grammar containing such rules can generate. In cognitive science, on the other hand, creativity is used more often in association with the ability to form new concepts, an ability that is linguistically reflected in productivity in the lexicon. But lexical productivity and recursiveness are quite different phenomena. And yet another sense in which language users may be said to be creative concerns not so much new lexical compounds or new sentential structures, but new uses of ‘old’ material. An obvious case in point is metaphor, but the formation of certain types of idiolect also qualifies. To call them all instances of ‘creativity’ is theoretically question begging since there is no *prima facie* reason to assume that we are dealing with the same phenomenon in each of these three cases. *A fortiori*, we should be reluctant to assume that properties of language and abilities of language users can be tied together as strictly as is often done. Such an association can not be a matter of theoretical stipulation, rather it should be the result of an independent, empirical investigation of the phenomena in question.

Yet in formal semantics, characteristically, such an investigation is lacking in the case of creativity and compositionality. A distinctive feature of formal semantics is that the two are considered to be a perfect match: creativity is compositionality.¹² But can creativity really be accounted for completely by giving a compositional theory of syntax and semantics? It should be noted this is a different point than the one that there may be other ways to account for creativity than through compositionality.¹³ That is quite a relevant issue to explore as well, but the present qualm works in the opposite direction. It asks

It is conceded by most philosophers of language, and recently by some linguists, that a satisfactory theory of meaning must give an account of how the meanings of sentences depend on the meanings of words. Unless such an account could be supplied for a particular language, it is argued, there would be no explaining the fact that we can learn the language: no explaining the fact that, on mastering a finite vocabulary and a finitely stated set of rules, we are prepared to produce and to understand any of a potential infinitude of sentences. I do not dispute these vague claims, in which I sense more than a kernel of truth. Instead I want to ask what it is for a theory to give an account of the kind adumbrated.

Notice that Davidson throws in yet another feature of language, viz., its learnability, and the corresponding ability of language users.

As an aside, note that both Montague and Davidson (explicitly and implicitly, respectively) appeal to Chomsky, which testifies to the enormous prestige of Chomsky and the revolutionary impact of his ideas at the time.

¹²This immediate association can be found in many discussions of compositionality. We already quoted Frege and Davidson above, here are two other representative quotes: [16, p. 281]: ‘...some version of it [= compositionality, G&S] would appear to be an essential part of any account of how meanings are assigned to novel sentences.’; and [4, p. 6]: ‘Both human thought and human language are, invariably, productive and systematic; and the only way that they could be is by being compositional.’

¹³Cf., [3] for a detailed defence of the thesis that we do not need compositionality in order to obtain an account of creativity.

whether the identification of creativity with compositionality is sound to begin with, i.e., whether it constitutes a possible explanation at all. The point is this. The empirical phenomena that we label with the phrase ‘creativity’ are real, and as such, constitute undisputed facts. But as we indicated above, we first need to investigate what they are facts about. Language? Competence? If the latter, what is the exact nature of this competence? Is it an individual ability, i.e., a property that can be assigned (and denied) of individuals? And can it be explained in this particular way, i.e., via identification with the kind of compositionality that we borrow from formal languages and then apply in our descriptions of natural languages? And these questions are of course related. If we proceed in the way that is characteristic for formal semantics, we locate creativity primarily in language and let competence mirror that. From which it follows that the latter needs to be individual. But whether that actually is the case, or not, is not something that should follow from a ‘mere’ methodological decision. Rather it is a substantial hypothesis that can be evaluated empirically: what competence is, should be established by observation, not by ideological fiat.

So let us take another look at compositionality as a presumed feature of natural languages. The idea that meanings can be assigned to complex expressions (sentences in particular) only in a compositional manner is obviously tied to the conception of language as an infinite object. ‘How many sentences of English are there?’ ‘Infinitely many.’ But if that is the case, then how can we characterise all of them and assign them a meaning in a finite manner? By using recursion in the syntax and relying on some form of compositionality in the semantics, the standard reasoning goes.

Notice, however, that this tends to gloss over an important issue: the nature of the infinity that is supposed to characterise natural languages. That the example set by formal languages plays a suggestive role here seems plausible. These are languages that have an infinite number of formulae and that (usually) have a compositional semantics. But note that formal languages are infinite by definition, literally: they are defined, designed, to be that way.¹⁴ But for a natural language this is not the case. Where it is a matter of decision whether we set up a particular formal language as an infinite object, surely for natural

¹⁴Another, more indirect source of inspiration may come from the way in which in elementary mathematics we convince someone of the infinity of the set of natural numbers, bears an obvious resemblance to the argument alluded to above. Here we explain that there can be no largest natural number, because we can always add 1 to any one candidate; and if there is no largest natural number there must be infinitely many. There are at least two things to bear in mind before we accept the transposition of this argument to the set of sentences of a given natural language. First of all, strictly speaking the argument is not compelling in the case of mathematics, as is witnessed by the existence of strictly finitistic approaches. (There are interesting parallels between the anthropological version of strict finitism à la Wittgenstein and Maddy in the philosophy of mathematics and the view on language that arises from our considerations. However, we can not explore them in this paper.) Secondly, it must be borne in mind that mathematics and linguistics deal with, prima facie at least, completely different types of objects. The existence and the properties of linguistic entities are empirical and contingent in ways that simply do not apply to mathematical entities. Hence, should one want to accept the argument, one needs an independent argument in favour of a platonic interpretation of linguistic entities. That seems not readily available.

languages the issue of its finiteness or infiniteness must be an empirical one, and hence of a completely different order? This much may be granted while at the same time one may consider the matter obvious: for is it not a correct empirical observation that for every sentence of English one can construct one that is longer, e.g., by using embedding or simple coordination? But what remains to be seen is how this bears on the issue of infinity. Obviously there are limitations on actual production and recognition of linguistic utterances, limitations that are not contingent in this sense that, although the actual values of the limits may vary, from individual to individual, their very existence is intrinsically related to the finitude of our cognitive capacities. What we individually will experience of a natural language is finite, as is what any one of us, past, present or future will. This is the observation of a humdrum fact, but it is important to realise, nevertheless, because it shows that the infinite nature we ascribe to a natural language as an object reflects a certain kind of abstraction. As such the abstraction may be a useful one, but it provides yet another indication that the idea of compositionality, based as it is on the presumed infinite nature of language, is a theoretical assumption, rather than an empirical fact.

And of course, what we commonly describe as the infinite nature of language can also be characterised in other ways, that do not commit us to the assumption that natural languages are infinite objects. What we, as competent users of a language, possess, is a certain ability, the result of an intricate set of relationships between a finite brain, a finite education, a finite social environment, etc., of which the actual application, although itself also finite (both individually and collectively) is not limited a priori. That is another, less metaphorical and more down to earth, way of describing the fact that, indeed, what utterances we will actually produce and interpret can not be determined in advance.

This is the first assumption that grounds the pivotal role of compositionality: that language is an infinite object, not just syntactically but also semantically, i.e., that there is a potentially infinite number of expressions with a potentially infinite number of meanings. The real strength of this assumption shows itself when it is combined with the second one, that competence is individual. This second assumption we already saw at work above, e.g., in the quotation from Frege in footnote 10, where he ‘identifies’ creativity of language with creativity of users. The idea that language is an individual asset, that whatever properties language has, they can be used, exploited by individual users (and hence are reflected in properties of these users), indeed is a cornerstone of formal semantics. Once these two assumptions, viz, that language is an infinite object, and that the individual is the locus of competence, are made, we face the task of reconciling in some way or other the infinity of language with the obvious finiteness of the individual brain (or mind, or whatever we prefer to locate competence in). But how to do that? How is a finite brain able to master an infinite language? That is the question of learnability, one form that the problem of creativity may take as we have seen above. Or: How is a finite mind able to understand (or to produce) an utterance of a sentence it has never before encountered? That is another form that the question may take. In every case we seem to be confronted with an apparent fundamental discrepancy between the infinity of languages

and the finitude of the individual. And it is precisely at that point, when we struggle with that insurmountable problem, that compositionality comes to the rescue. It allows us to come up with a finite representation of this infinite object, one that we can consistently assume to be mastered and used by a finite individual.

If we combine the result with yet another deeply rooted conviction, viz., that the mental is the primary, if not the exclusively relevant aspect of the individual, what emerges from these observations is a very specific, and philosophically quite ‘loaded’, view on what constitutes the domain of linguistic theory (i.e., of syntax, semantics and their interface): it basically consists of an individual, mental ability to autonomously produce and interpret a potential infinite set of linguistic utterances. And this picture of language, meaning and competence is characteristic not just for formal semantics, but for many other twentieth century paradigms as well. Alternative views notwithstanding, it is an important characteristic also of linguistics in the Chomskyan tradition, broadly conceived, and of the dominant approach in cognitive science. Characteristic for this view, of which the historical roots go back to the beginnings of modern philosophy, is the fundamental assumption that the use of language is an expression of thought, that what we say and how we say it reflects what we think and how we think it.

It is against this background that compositional semantical analyses are given and that the various arguments concerning the nature of compositionality, its formalisation and its feasibility, take place. In that sense compositionality is an internal issue, something that arises and plays the role that it does in virtue of these fundamental assumptions. It serves primarily an internal function, since it is only given the assumptions identified above and given the resulting view on the organisation of grammars, that doing things compositionally becomes both a requirement and at the same time a possibility. It is these theoretical views on language and mind that make compositionality possible and it is in this sense that compositionality itself is a theoretical principle too.

Of course, it is one thing to note that the assumptions behind compositionality (too) often go unnoticed, it is another to spell out their consequences. So now that we noticed them, what conclusions can be drawn? After all formal semantics, and modern linguistics in general, exemplifies an important and successful research strategy. Is there any reason to think that reflection on these assumption will change that? The answer to that question is probably ‘No, at least not any time soon.’ But issues concerning compositionality remain, as is exemplified by the literature referenced earlier. In addition, further consideration of these underlying assumptions may well shed some more light on the kind of discipline that formal semantics is, and thereby bring into focus what we can and can not expect semantics to contribute to other scientific endeavours. In what follows we will briefly look at two instances. The first concerns the relevance of compositionality for the relationship between language and thought as conceived along the lines sketched above. The second is concerned with the challenge provided by alternative views, in particular those that question the special status of the individual mental.

3 The Thought Gang

As we indicated above, the idea of an intimate relationship between language and thought is one of the cornerstones of dominant paradigms in both linguistics and cognitive science. It is also characteristic for an entire tradition in philosophy. One form the idea takes is, as we also saw earlier, the identification of compositionality of meaning and creativity of individual language users. From a cognitivist point of view this identification becomes problematic once doubts arise concerning compositionality. In a recent paper¹⁵ Jerry Fodor deals with this issue and gives the discussion a positive twist by turning it into an argument in the long standing debate about the priority of language and thought. His discussion identifies the following question as forming the kernel of that debate: ‘Is it thoughts or sentences that “have content in the first instance”?’ he asks.¹⁶ What is interesting for our present concern is that according to Fodor compositionality will decide this issue:

As between the two, at least one of thought and language must be compositional. If *only* one of them is, then that’s the one that has content in the first place.¹⁷

As usual, Fodor’s position here is crystal clear: compositionality is imperative for having underived content,¹⁸ natural languages are not compositional, hence thought is prior to language when it comes to having content. That linguistic meaning can not be described in a compositional fashion, Fodor concludes on the basis of an analysis of the way in which the meanings of definite descriptions contribute to the meanings of sentences in which they occur.¹⁹

In what follows we will not be concerned with what for Fodor is one of the main topics he addresses in the paper, viz., the impossibility of having what he calls ‘an epistemically based semantics’, i.e., the untenability of ‘the idea that semantics might underwrite justification.’²⁰ That supposition, that epistemological issues concerning, e.g., justification, could be solved via an analysis of the semantics of the expressions used in epistemic claims, Fodor argues, is ‘mistaken root and branch’²¹. When it comes to justification of epistemic claims,

¹⁵Cf., [4].

¹⁶Ibid., p.2.

¹⁷Ibid., p. 7.

¹⁸‘... the thesis that underived content is compositional is not negotiable’, *ibid.*, p. 6.

¹⁹ It should be noted that Fodor talks about compositionality without spelling out either ‘structure’ or ‘meaning’, since he assumes this is not needed for his argument (*ibid.*, p. 6). According to Szabó (cf., [21]) that puts him in the methodological camp. However, Fodor concludes that as a matter of fact language is not compositional, a claim that seems indeed difficult to assess without some specification of the relevant notions of ‘structure’ and ‘meaning’. Also note that Fodor uses compositionality in a strict sense: only properties of constituents determine properties of the whole and all of them do (*ibid.*, p. 9). That actually does rely on a spelling out of structure and content, that Fodor thinks he does not need to give. However, this strict interpretation is crucial for his argument. Cf., also [16, p. 291 ff.], on Hintikka’s ‘determinacy thesis’ and the challenges that thesis faces.

²⁰[4, p. 4].

²¹*Ibid.*, p. 10

he says, it is the relationship between the content of those claims and the way the world is that matters, not the meanings of the sentences expressing those claims. These are issues we can not go into here. Instead we will concentrate on the assumptions underlying Fodor's reasoning, since they provide concrete illustrations of the general assumption we identified in the previous section.

A first thing to note is that according to Fodor compositionality is a property that, at least in principle, applies equally to language and to thought. Indeed, if compositionality were not a possible property of language, the whole question need not be raised in the first place. This in itself seems a strong assumption, one that ties basic features of language and of thought to each other, but one that remains unargued in Fodor's paper.²² In fact, Fodor's argument centres around the possibility of our being able to compare in some way the content of a thought and the meaning of a sentence that is used to express it. That comparison reveals that the two need not be in line, that, as Fodor puts it, 'language is strikingly elliptical and inexplicit about the thoughts that it expresses'²³ and 'as a simple matter of fact, in the general case, sentences are remarkably *inexplicit* with respect to how the thoughts they express are put together.'²⁴ This then leads Fodor to conclude quite straightforwardly that language is not compositional:

Either the content of the thought is different from the content of the sentence that expresses it, or the sentence isn't compositional. I take it that the first disjunct is preposterous; so I take it that the second disjunct must be true.²⁵

Of course the first disjunct is preposterous only if you are willing to assume that thoughts are the kinds of entities that get expressed by sentences, i.e., that thoughts have an existence of their own, in a realm of their own, independent of whether, and if so how, they are expressed. But that assumption is dubious, and once we give it up, the first disjunct is no longer preposterous. Rather, any claim that compares contents of thoughts with contents of sentences simply stops making sense. So we see that it is only given a specific assumption about the nature of thought and its relation to language that this argument carries any weight. If we fore-go the idea that thoughts are an individual commodity waiting to be expressed by the one who has them in whatever language that individual has at his or her disposal, then the entire problematic, including the conclusions about compositionality, vanishes.²⁶

²²That Fodor regards compositionality as a central issue in cognitive science as such is clear. See for example, the objections he and Pylyshyn (in [5]) raised against connectionism. Their arguments have been criticised by a number of authors, also on the point of the role of compositionality. See [1] for an assessment.

²³[4, p. 11].

²⁴Ibid., p. 12.

²⁵Ibid., p. 12.

²⁶Note also that 'productivity' (along with systematicity) is what according to Fodor requires compositionality: 'Both human thought and human language are, invariably, productive and systematic; and *the only way that they could be* is by being compositional.' (p. 6; our emphasis). What is at stake here is the identification of compositionality and competence. Fischer ([3]) does not agree, in as much as he rejects the 'tacit knowledge' claim that it leads to.

Let us now turn to the arguments that Fodor gives to substantiate his claim that language is, as he phrases it, ‘strikingly elliptical and inexplicit’. He briefly discusses two kinds of cases.

The first type of phenomenon concerns the underdetermination of what one intends to communicate by the meaning of the actual utterance. Basically, what Fodor has in mind here is the use of indexical expressions as a way of complying with Gricean requirements concerning communicative efficiency. As an argument for his thesis this looks suspiciously like ‘begging the question’:²⁷ only if we assume that the meaning is not strictly dependent on the context of utterance can we say that the meaning of ‘It is four o’clock’ underdetermines our intention (to say that it is four in the afternoon, or four at night, respectively). So we assume rather than show that linguistic meaning is independent of context.²⁸

The second phenomenon is meant to show that meaning when compositionally construed is not what is conventionally expressed. Fodor’s example is that of definite descriptions. Russell, he states, showed us how to construe sentences with definite descriptions compositionally²⁹, but in reality the Russellian presuppositions of existence and especially uniqueness are (almost) never satisfied. That claim is true, but hardly sufficient to make the point. Two things need to be added, viz., that no alternative, compositional analyses are feasible, and that an adequate semantic account necessarily should reflect existence and uniqueness of some analogous element in thought.

The first point is addressed by Fodor: he summarily dismisses attempts to give an alternative compositional semantics, referring to attempts to analyse descriptions as demonstratives and asking: ‘Why on earth any remotely sane language would use as a demonstrative an expressions with the syntactic structure of a quantifier is, however, not explained.’³⁰ But this will hardly do. Note that Fodor simply assumes, but does not argue, and how could he, that we have theory-independent ways of establishing what a demonstrative or a quantificational expression is. But of course we do not, so the question he asks is merely a rhetorical way of stating his point, not an argument for it.

The second point is not argued by Fodor explicitly, but seems simply taken for granted. In this respect it is interesting to note that Fodor is not prepared to make the semantics of thought equally a matter of fact: ‘a thought can’t be inexplicit with respect to its own content . . . a mental representation is ipso facto compositional with respect to the content that a correct semantics would assign to it.’³¹ Clearly this is an a priori claim about a priori properties of thought.

²⁷Cf., also what was remarked above, in footnote 19.

²⁸And in doing so we block the possibility of an externalist, and in the end a social, theory of meaning; cf., section 4.

²⁹Which is an odd claim given the fact that Russell could only give a contextual definition, and it was only with the advent of the use of higher-order logic and lambda-abstraction that a really compositional account could be given.

³⁰[4, p. 13].

³¹Ibid., p. 14. Cf., however Fodor’s footnote 2 (p. 7): ‘. . . ditto for thoughts assuming that the representational theories [*sic*] of mind is true.’ This is remarkable, since it suggests that a failure of compositionality actually falsifies representationalism. Does this mean that despite

Unfortunately, Fodor never says how this claim should be substantiated.

The upshot of his considerations for Fodor is this. Semantics is not compositional, yet compositionality is not negotiable. So, language does not have a semantics. A surprising conclusion. What then is it that semantics studies? According to Fodor, ‘it’s the study of the (compositional) semantics of thought’.³²

But are these conclusions forced upon us, and thereby also the assumptions that go into Fodor’s argument for them? No. That languages are not (completely) compositional may also lead one to explore a different account of their systematicity: not as elliptic, and inexplicit expressions of perfectly transparent and a priori structured independent thoughts, but as tools that we have acquired and are using in exchanges with others in a linguistic community that depend essentially on the context in which they are put to use.³³ Analysis of definite descriptions and their kin in the context of dynamic semantics, we venture, can be read as elements of such an account.³⁴

Given that possibility, it is to be deplored that Fodor does not feel he needs to argue the assumptions from which the positions he discusses, start. And that is not because he is not aware of the existence of alternatives:

There are, of course, other positions in logical space; it could be that the content of thought and language both derive from some *third thing*; or that neither thought nor language has content; . . . But none of these options strikes me as attractive and I won’t discuss them in what follows.³⁵

And that is a pity, for we surely could have learnt a lot if Fodor had discussed them. What is striking is Fodor’s justification: whether an option is attractive or not seems rather a dubious criterion for deciding whether to take it into account; surely, what should count is whether it makes sense, whether it provides insight, perhaps even (dare we say it?) whether it is true.

Actually, from another tradition there is one such obvious candidate option, viz., the view that both language and thought have content because and in so

all the ‘compositionality is not negotiable’ talk of Fodor’s there is a real issue here?

³²Ibid., p. 13.

³³Space does not permit extensive discussion here, but Wittgenstein’s observations in *Remarks on the Philosophy of Psychology*, part I, section 40 ff. ([25]) may provide a starting point for a fundamental criticism of this aspect of Fodor’s position. There Wittgenstein discusses the conceptual confusion that arises if we try to make sense of the prima facie plausible idea that in a sentence such as ‘A. Schweizer ist kein Schweizer’ we ‘mean’ the first occurrence as a name and the second as a noun. This picture of identifiable ‘though elements’ corresponding to ‘meaning bodies’ is muddled, and it is compositionality along with the kind of conception of the relationship between language and thought that Fodor seems to adopt, that leads to this confusion. Cf., also [24, §551–§566].

³⁴That compositionality is such an issue in cognitive science is directly related to the question of representationalism, that is so central for the symbolic paradigm. Dynamic semantics presents a conscious attempt to do away with representations. In doing so, it also moves away from the referential word–world approach to meaning in favour of a communicative perspective: the individualistic subject–world orientation is being replaced by the intersubjective (social) subject–subject orientation. Cf., [9]; see [18] for discussion.

³⁵[4, p. 2].

far as they are related to the public practices of a (linguistic) community. Let us briefly look at that.

4 Compositionality in practice

That language and meaning are not strictly individual assets but belong to a community, is probably a claim that strikes many as patently false, if not absurd. After all, it seems that if I use language, in speaking or writing, listening or reading, I am able to do so precisely because I, as a individual, possess a sufficient grasp of the language that I use. Sure, I may need to consult a dictionary or a fellow language user every now and then, but basically I am competent on my own. However, a case can be made that language and competence need to be construed as intrinsically social.

A first source of considerations derives from what Putnam aptly dubbed ‘the division of linguistic labour’,³⁶ viz., the fact that in many cases the properties connected with certain words that determine their reference are not known by all members of the linguistic community, but are privy to a small set of ‘experts’. We all know how to use a host of terms referring to chemical substances, biological species, physical phenomena, astronomical entities, historical and literary figures and events, and so on and so forth, without being able to specify the relevant semantic features of these terms that determine their reference. In as much as meaning determines reference, then, this aspect of it is not one that we may assume all users of the language have a grasp of. The effects of this social nature of meaning stretches beyond the determination of reference. It also affects, for example, lexical productivity: it make take ‘expert knowledge’ to determine whether a particular verb can be turned into a dispositional adjective (depending on its material constitution), whether a certain verb is or is not transitive (or di-transitive), or does or does not go with certain PP’s. In these cases compositionality is ‘a communal effort’, i.e., these are examples of semantic knowledge, in the sense of knowledge that plays a role in composing (producing) and decomposing (interpreting) an utterance, that is essentially communal knowledge, not individual.³⁷

The alternative picture takes seriously that people become language users only as members of a community, that outside a community there is no linguistic competence for an individual human being to develop. These are simple facts, that are, however, not always taken on face value. In the Chomskyan tradition

³⁶Cf., [17].

³⁷This externalism with respect to meaning extends, it can be argued, to mental content. On the assumption that meanings are somehow related to mental content, the contents corresponding to these shared terms are shared as well: mental content is, at least partly, determined externally. So-called ‘two-factor’ theories (Block, Recanati) are attempts to deal with external determination of reference that aim to preserve compositionality of narrow content and the identification of narrow content with mental content. Cf., [19] for discussion. Since our primary focus is meaning, not mental content, we feel we sidestep the issue of mental content here. What is relevant to note is that in many of these analyses the individualistic bias is still present to such an extent that competence is regarded as basically individual.

one still insists that the capacity to develop language is individual, and that this individual capacity only needs ‘nourishment’ from the linguistic community to develop itself into full grown linguistic competence. Like a plant that contains the entire blueprint of its growth, flowering and fruit bearing inside itself, and merely needs nutrients, water and light from its environment to realise the blueprint. The problem with this picture is that it isolates linguistic competence from various other cognitive and emotional features that characterise human beings. And that it fails to recognise to what extent the individual subject as such is shaped both in terms of its capacities as their employment by the community it is part of. It is external factors, both synchronically as well as diachronically, that help to shape what individual competence is.

The question is not whether or not an individual can be said to be competent: of course that can be said (and denied). Rather the issue is whether what saying that an individual is competent (or not) means, itself can be explained purely in terms of individual properties, i.e., in terms of the individual’s physiological and/or psychological make-up, without any essential reference to other factors: other language users, the linguistic community, historical factors. That, it is claimed, is not the case: the idea of competence being individual in the sense that it could, in principle if not in fact, depend solely on strictly individual properties, so that the case of a single individual possessing semantic competence in absence of any other either synchronically or diachronically, is not excluded, that idea is conceptually incoherent (and not just contingently impossible). And that implies that the ‘burden of creativity’ need not be carried by the individual ability alone.

Note that the switch from individual to social does not eliminate the ‘problem’ as such: a linguistic community is after all nothing but a finite number of finite individuals. So we need an additional argument why we need not worry (not too much at least). And that may be this: that the limits of the community ipso facto constitute the limits of language in this sense that there is no ‘outside’ authority that would be able to come up with an utterance that is novel to the community as a whole. To be precise: if a new utterance presents itself, it is always produced by someone inside the community, hence not totally new. So the really troublesome situation that arises with a single individual as locus of competence: that there is this single finite competent speaker that is confronted by whatever ‘language’ wants to throw at it, does not arise here.

Perhaps the following provides a nice illustration of how a social conception of meaning might reflect on the issue of compositionality. We know from Wittgenstein’s work that self-attributions of knowledge in combination with self-ascriptions of such predicates as ‘being in pain’, are ‘out’.³⁸ The reason is as simple as profound. The concept of knowledge is, Wittgenstein argues, intrinsically connected with such concepts such as doubt, ground, (positive and negative) evidence. So, if we say ‘I know that p ’, that utterance signals the absence of doubt, our confidence that we are able to give convincing evidence for p being the case, and so on. Wittgenstein also notes that to say that this is

³⁸Cf., [24, §§ 246, 408].

the case makes sense only in cases where, in principle, the ‘counterparts’ of these notions, i.e., the existence of doubt, the availability of evidence to the contrary, etc., make sense as well. One can only be confident when the possibility of doubt is not excluded, and one can only bring forth positive evidence in cases where negative evidence is imaginable. (Of course, the doubt need not be actual and the negative evidence need not actually exist.) All this, Wittgenstein notes, fails in the case of ‘I know that I am in pain’. For what could make me doubt that? What evidence could I give? What could convince me that I am wrong? Under normal circumstances these questions do not make sense, and hence, neither does the utterance itself, Wittgenstein concludes. At best ‘I know that I am in pain’ is an emphatic way of stating that I am in pain, but unlike its second and third person counterparts (‘You know / he knows that I am in pain’) it is not a claim about the epistemic state of an individual.

What is interesting for our concerns with compositionality is that this is neither a feature of the concept of knowledge, nor of that of pain, or that of the first person, as such. It arises out of their combination and interaction. So what we have here is truly a matter of compositionality beyond the domain of concepts and concept formation: this is ‘compositionality in practice’. Important also is the relationship with the subject: this is one place where the social constitution of the subject (which includes the subject to which we traditionally ascribe semantic competence!) becomes apparent. Interesting also is that ‘propositional attitude ascriptions’, of which this is a particular case, are a field at which many semantic theories perform rather poorly. In particular when we focus on the meanings of those attitude expressions that are closely connected with meaning and semantics, such as ‘to intend’, ‘to mean’, ‘to understand’, ‘to refer’, ‘to describe’, and take seriously that the meanings of these very expressions originate, not in soliloquy or through introspection, but in conversation and interactive learning, it becomes clear that the capacities to which these terms refer are social through and through.

This points to the following, more general conclusion. In many cases ‘externalism’ with respect to concepts can be accommodated to the individualistic bias of semantics by the assumption that for every external determinant of meaning there exists an element in the individual mental make-up and that it is the latter, not the former, that semantics deals with. For example, assuming with Putnam that the extension of ‘water’ is determined externally, viz., by the essential nature of the stuff itself, we can nevertheless hold that the meaning of ‘water’ is individual in that it provides a pointer to that external determinant, but not the external determinant itself. However, such an escape hatch is not provided in the case of first versus third person asymmetries, such as in the case of ‘pain’. The gulf being epistemological, not semantic, can not be bridged in such a way, and therefore has to be acknowledged and accounted for in the semantics itself. And that means that a semantics of pain-ascriptions can not be had purely compositionally.

5 The discipline of semanticists

What are the consequences of the above for semantics, specifically for its disciplinary status? Elsewhere³⁹ we have argued that it is not completely obvious that semantics is an empirical discipline. The phenomena that semantics deals with are too much of a theoretical nature to allow for such a straightforward claim. Of course, empirical considerations do constrain semantic descriptions and analyses, and inform semantic theories. But the very nature of meaning, standing as it does at the cross roads of a host of different kinds of empirical phenomena — social, cognitive, formal, historical, biological — defy any simple and straightforward determination of the enterprise that attempts to characterise it. And the considerations concerning compositionality presented in this paper, we feel, underscore this.

Let us give one illustration how these considerations may affect what we actually do in semantics. One other reason why compositionality is such an important concern within semantics is that it allows us to assume just those levels of representation that we assume we have access to, viz., form (syntax) and meaning (semantics). That is to say, compositionality allows us to avoid positing ‘in between’ levels, such as that of LF, at the sentence level, or that of DRSS, at the discourse level.⁴⁰ But to what extent and in what way form and meaning are indeed accessible is a moot point, of course. As far as form is concerned, typically structuralist considerations do show that language is structured. However, what that structure is, whether it can be captured in categories or in dependencies, whether it is flat versus deep, etc., is beyond direct access by such means, and hence highly theory-dependent. This is not to say that there may not be additional empirical evidence that is relevant, but the fact remains that no clearcut empirical considerations allow us to choose between completely different, but descriptively equally successful rivals, such as categorial grammar, minimalism, HPSG and LFG like grammars, dependency grammars, etc. The formal counterpart of this situation is, of course, that weakly equivalent grammars exist for one and the same language. As for meaning, the usual ‘empirical’ access here is supposed to be furnished by introspection. But that is not really an answer. As far as meaning is concerned, we have to make do, at least initially, with behaviour (both verbal and non-verbal), and may postulate the existence of meanings only as invariants over behaviour. These invariants, then, can be studied reasonably independently from the behaviour and this independent study then is mistakenly identified with intuitive access to a realm of independent entities.

What does this imply for the way in which we describe language? We can describe language as if it were compositional and infinite, and competence as if it were individual, as long as we realise the ‘as if’ we are working with. To put it differently, we can, if we wish, if it is useful, practical to do, if it leads to

³⁹Cf., [20].

⁴⁰Evidence for this role of the compositionality assumption comes both from ‘negative’ as well as from ‘positive’ uses of compositionality. Cf., e.g., the use of compositionality considerations used in favour of LF in [10], and those used against DRSS in [8].

formally explicit and elegant theories, describe language compositionally. But the resulting system is not one which we can regard as a characterisation of the object of individual competence. Compositionality is a formal property, an attractive one and often a useful one, of an entity that itself is an abstraction. There is nothing wrong with that, but it does raise new questions concerning the relationship between this type of semantics (linguistics) and actual language use (including cognitive aspects).

As we remarked above, meaning is a very complicated notion. Sometimes it is claimed that ‘meaning’ is a theoretical concept, just like ‘mass’ in physics, and that therefore in semantics just as in physics we are allowed to abstract from various features of the empirical phenomena that we are dealing with. That, however, will not do. First of all, it is not clear that ‘meaning’ is a theoretical concept to begin with: it is something that in everyday discourse we refer to, talk about, discuss, a concept that we are directly involved with in a myriad ways.⁴¹ Secondly, and more importantly, unlike in physics, what is lacking in semantics is a proper operationalisation of the concept in question, one that is independent of the theoretical framework. And do note that an appeal to intuitions (about meanings, entailments, etc.) does not provide that. This suggests that we require not only that the formal concepts that we use should be adequate with respect to the informal concepts that are being analysed, but also that it should be possible to translate them back to these informal concepts. This is directly related to the nature of the phenomena we are dealing with: what we model, i.e., should be modelling, is what meaning is for those that use it. With this strong constraint in place we pinpoint the usefulness of formalisation to exactly the right issues: precision, clarity, computability, etc.

The bottom line is that of course we can still do compositional semantics, but we should realise that in doing so we start from certain assumptions that, although convenient, are not accurate. For semantics as such, this does not change much. But it does affect, at least in principle, some of the connections that are usually assumed to exist between semantics and other disciplines. One notorious example is the relationship with cognitive science. If the considerations we outlined in this paper are on the right track, we can no longer work on the traditional assumption that since language and thought mirror each other, semantics studies a cognitively relevant entity. At least, not when we stick to our compositional methodology. So, it seems that there are consequences after all, but these do not so much concern what semanticists *do* but rather what semantics *is*.

References

- [1] Keith Butler. Connectionism, classical cognitivism and the relation between cognitive and implementational levels of analysis. *Philosophical Psychology*, 6(3):321–33, 1993.

⁴¹Cf., [20, §3] for more discussions.

- [2] Donald Davidson. Truth and meaning. *Synthese*, 17:304–23, 1967.
- [3] Eugen Fischer. Bogus mystery about linguistic competence. *Synthese*, 135:49–75, 2003.
- [4] Jerry A. Fodor. Language, thought and compositionality. *Mind and Language*, 16(1):1–15, 2001.
- [5] Jerry A. Fodor and Zenon Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28:3–71, 1988.
- [6] Gottlob Frege. Gedankengefüge. *Beiträge zur Philosophie des deutschen Idealismus*, 3. Band:36–51, 1923–1926. Reprinted in Geach (ed), 1977.
- [7] Peter Geach, editor. *Logical Investigations*. Blackwell, Oxford, 1977.
- [8] Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14(1):39–100, 1991.
- [9] Jeroen Groenendijk and Martin Stokhof. Meaning in motion. In *Reference and Anaphoric Relations*, pages 47–76. Kluwer, Dordrecht, 2000.
- [10] Irene Heim and Angelika Kratzer. *Semantics in Generative Grammar*. Blackwell, Oxford, 1998.
- [11] Theo M. V. Janssen. Frege, contextuality and compositionality. *Journal of Logic, Language and Information*, 10:115–36, 2001.
- [12] Theo M.V. Janssen. *Foundations and Applications of Montague Grammar*. PhD thesis, Department of Mathematics, University of Amsterdam, 1983.
- [13] Theo M.V. Janssen. Compositionality. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Linguistics*, pages 417–74. Elsevier/MIT Press, Amsterdam/Cambridge Mass., 1997.
- [14] A. Kazmi and J. Pelletier. Is compositionality vacuous? *Linguistics and Philosophy*, 21:629–33, 1998.
- [15] R. Montague. Universal grammar. *Theoria*, 36:373–98, 1970.
- [16] Barbara Partee. Compositionality. In Fred Landman and Frank Veltman, editors, *Varieties of Formal Semantics*, pages 281–311. Foris, Dordrecht, 1984.
- [17] Hilary Putnam. The meaning of ‘meaning’. In *Mind, Language, and Reality*, pages 215–71. Cambridge University Press, Cambridge, 1975.
- [18] Manuel Rebuschi. Towards a neutral semantics. *Linguistics and Philosophy*, to appear.
- [19] Philip Robbins. How to blunt the sword of compositionality. *Noûs*, 36(2):313–34, 2002.

- [20] Martin Stokhof. Meaning, interpretation and semantics. In D. Barker-Plummer, D. Beaver, J. van Benthem, and P. Scotto di Luzio, editors, *Words, Proofs and Diagrams*, pages 217–40. CSLI, Stanford, 2002.
- [21] Zoltán Gendler Szabó. Compositionality as supervenience. *Linguistics and Philosophy*, 23:475–505, 2000.
- [22] D. Westerståhl. On mathematical proofs of the vacuity of compositionality. *Linguistics and Philosophy*, 21:635–43, 1998.
- [23] D. Westerståhl. On the compositionality of idioms. In D. Barker-Plummer, D. Beaver, J. van Benthem, and P. Scotto di Luzio, editors, *Words, Proofs and Diagrams*, pages 241–71. CSLI, Stanford, 2002.
- [24] Ludwig Wittgenstein. *Philosophical Investigations*. Blackwell, Oxford, 2nd edition, 1958.
- [25] Ludwig Wittgenstein. *Remarks on the Philosophy of Psychology*. Blackwell, Oxford, 1980.
- [26] W Zadrozny. From compositional to systematic semantics. *Linguistics and Philosophy*, 17:329–42, 1994.