

Comments-to-comments

MARTIN STOKHOF and MICHIEL VAN LAMBALGEN

Preliminary

Let us start by thanking the commentators for their many thoughtful and critical comments. We are happy that the main thesis we set out to defend in the paper has given rise to such interesting observations and (counter)arguments. Most of them are critical of the analysis of the role of idealisation in modern linguistics, and have given us much food for further thought. Unfortunately, not every comment made by the commentators can be discussed separately. In what follows we want to keep the discussion as systematic as possible, and therefore we will group various comments under three headings, and address these topics as systematic issues, but of course we will refer to the arguments of the individual authors whenever this is relevant.

The topics we want to discuss are the following. First of all, we want to be more specific about the scope of the analysis given in the paper: to what paradigm and/or part of modern linguistics does our analysis apply? (That is, if to any at all . . .). Secondly, we want to address various comments and criticisms regarding the nature of idealisation and the applicability of this concept to linguistics. Then we will address the question whether there is indeed a difference at this point between linguistics and the sciences. Fourthly, we want to further elucidate our use of the label ‘ideological’ in connection with how certain idealisations are motivated, and we end with a remark that addresses the actuality of our analysis.

1. The scope of the analysis

As it states in its very title, the paper analyses the role of idealisations and abstractions as part of ‘the construction of modern linguistics’. Several

commentators have taken issue with this,¹ and from a certain perspective we think they have a point. If one looks at it from a present-day perspective, and takes ‘modern linguistics’ to denote linguistics as we have it today, then, indeed, our analysis targets only a specific, and arguably less and less influential part of that. For the larger part (but cf. below for some qualifications), the paper addresses issues that concern the generative tradition in linguistics, and in the paper the examples are taken only from work in that paradigm. However, as we had hoped we had made clear in the introduction of the paper, when we used the phrase ‘modern linguistics’ we meant it as referring to the great change in linguistics that occurred in the fifties and sixties of the last century, when linguistics became a fashionable discipline, and, in the eyes of many, both an exciting and a respectable branch of science. That development, for better or worse and justifiably or not, is for the larger part the development of the generative tradition as framed by Chomsky and his associates. It is certainly the case that, in terms of numbers of researchers actively working in that paradigm, the generative tradition has been in decline for some time now. It can also be argued that conceptually there always have been severe problems with the entire approach (indeed, this is one of the main contentions of the paper), and that right from the early days there have always been critics that have pointed out its various flaws. But that does not diminish the fact that in its heyday generativism was intellectually dominant, and as such has shaped the ‘persona’ of modern linguistics to a large extent.

That being said, although we did not explicitly argue this to any extent in the paper, we are of the opinion that some of the basic constructions that originated with the generative tradition in syntax continue to be operative also in other paradigms and other linguistic disciplines. And this should not come as a surprise of course. For one thing, many such paradigms are explicitly developed as alternatives to the generative one, and no alternative does not share some basic assumptions and concepts with the original it wants to supersede. Also, the diversity that is characteristic of present-day linguistics is the result of the extension of methods and concepts that were originally developed in syntax to

¹ Cf., e.g., Newmeyer and Stabler. Our use of ‘modern linguistics’ was correctly diagnosed by Haspelmath.

other branches of linguistics, such as semantics and pragmatics. Thus we think it is fairly obvious that despite the many and important differences that exist between various alternative paradigms, these also share a number of core assumptions and concepts with the generative tradition. And in as much as these assumptions and concepts are instances of what we have labelled ‘idealisation’, that means that our analysis in effect extends to these alternative paradigms as well, and thus that, appearances notwithstanding, it does range over more than just the generative tradition in ‘modern linguistics’ even if we take the latter term to refer to present-day linguistics.

By way of illustration let us briefly mention two such features that seem to have survived the demise (if such it is) of the generative paradigm in syntax and that are, knowingly or unknowingly, perpetuated in different frameworks and linguistic disciplines. They are the infinite nature of language and the notion of competence.

That language is an infinite object, i.e., that potentially each natural language has an infinite number of expressions, is an assumption that is widely shared in modern linguistics. One may find it in the generative tradition in syntax, in formal semantics, in optimality theory, and so on. How specific an assumption is this, however? Stabler notes that it is also a common assumption in computational models of a wide variety of phenomena, including human problem solving. But it seems the assumption does not always function in the same way. Take the case of action planning. Although a computational model might actually subscribe to it for the sake of convenience, there is nothing in the phenomenon as such that forces upon us the assumption that humans actually have an infinite number of plans or action options at their disposal when engaged in action planning and decision making. Rather, the more natural assumption is that, given a goal, a plan is computed on demand, which is in fact more like a constraint-based grammar than like a generative grammar.

So, it seems the assumption that language is an infinite object is rather specific. In the paper we referred to Tomalin (2006), where it is argued that the origins of Chomsky’s use of this concept is mathematical logic and computation theory as it was developed in the first half of the last century. It is true that the concept of natural language as an infinite structure is widely embraced in linguistics, but the question is, of course, whether the fact that it is so commonly embraced adds to the plausibility of this assumption. An answer to that question requires an in-depth investigation into the nature of the arguments

that are given for it in the literature. We cannot go into great detail here,² but perhaps the following suffices to illustrate our point.

One standard way of defending the assumption that we can find in linguistics is by a simple argument, which appears to be modelled after a similar argument in mathematics that aims to establish that the set of natural numbers is infinite. Just as for any purported largest natural number we can construct a larger one by adding 1 to it, the argument goes, we can show that there is no longest sentence, since we can always embed a purported candidate in another one. Consider the following characteristic instance of this argument, taken from a paper by Howard Lasnik (Lasnik, 2005, p. 64):

One of the most fundamental properties of human language is its infinitude: there is no upper bound on the length of sentences, hence no upper bound on the number of sentences. In all human languages, ever longer sentences can be constructed by embedding one sentence inside another, as illustrated by:

- a. *Mary reads books.*
- b. *John thinks Mary reads books.*
- c. *Susan said John thinks Mary reads books.*

The argument depends on the assumption that some rules in the syntax of a natural language are recursive. In fact, recursion is widely considered to be a characteristic feature of human language (as opposed to the communication systems of other animals).³ However, it is quite obvious that no observation of actual linguistic material really forces that assumption on us. For example, Karlsson (2002) reports on the basis of a study of large corpora of both spoken and written language (English and Finnish) that the ‘usus maximum’ of finite right embedding is level 4. On the semantic side, it is known from studies on embedded theory of mind (Rutherford, 2004) that humans process embeddings up to level 2 faultlessly, but already start to make mistakes with embeddings of level 3.⁴

² But, fortunately, we can refer the reader to the discussion in Pullum & Scholz (2010), the conclusion of which is that none of the standard arguments that are around actually necessitate the assumption that language is infinite, or has to be treated as such.

³ Though it should be noted that recursion is neither sufficient nor necessary for infiniteness. For an example of a recursive grammar that characterises a finite language, cf., Pullum & Scholz (2010).

⁴ Cf., also Fitz (2009) for a study of artificial learning models that are able to deal with embedding constructions of finite level without depending on recursive rules.

Now one often reads that even if actual linguistic material may not strictly validate the assumption of the recursive nature of natural language, it is still the most elegant way to account for the existence of various nesting and embedding constructions. We do not want to take issue with the appeal to elegance, but merely want to point out that such an account, elegant or not, makes predictions that obviously are not borne out by the facts, and thus saddles us with the additional task of finding an explanation of the discrepancy. This supports the claim made in the paper that we are dealing with an idealisation here, i.e., with a type of construction that by shifting ontologies creates additional epistemological burdens.

To put the same matter slightly differently, the association between recursion and human language, not being based on any observable facts, but merely following a conceptual argument borrowed from an entirely different domain, viz., that of mathematics,⁵ is conceptual rather than factual itself. And this then also holds for the assumed infinite nature of language as such: it is a characteristic feature of a concept rather than an observable property of a natural phenomenon.

What is relevant to observe is that the line of reasoning briefly indicated appears 'across the board' in present-day linguistics: not just in the generative paradigm, but also in other frameworks. Three examples should suffice. The first comes from Pinker and Jackendoff, who, in a highly critical review of Hauser et al. (2002), accept as a given that 'human languages are recursive', and hence infinite (Pinker & Jackendoff, 2005). What they disagree with is the claim that recursivity all by itself is characteristic of human language.

The second example comes from optimality theory. In Tesar et al. (1999, p. 313) the set of candidate forms is recognised as infinite in nature, due to recursive mechanisms:

The infinite size of the candidate set results from the possibility of candidates containing material not found in the corresponding input; epenthesis (insertion) in the CV syllable theory is an example. Once the possibility of such insertions is granted (a rather uncontroversial position), there isn't any principled basis for placing any absolute limit on the number of insertions a candidate description may contain.

A different paradigm, to be sure, but one that accepts the same assumption.

⁵ It is for that reason that our qualms with regard to the argument as applied to natural language do not carry over to the domain of mathematics.

As a third example, note that in formal semantics the infinite nature of language is reflected in the use of formal languages as models for natural language. The well-known claim in the opening passage of Montague's 'Universal Grammar' (Montague, 1970), where Montague states: 'There is in my opinion no important theoretical difference between natural languages and the artificial languages of logicians', provides a clear illustration: both formal languages and natural languages are assumed to be infinite sets of well-formed expressions.

The assumption that language is an infinite structure is thus definitely not one that is confined to just generative syntax. On the contrary, as these examples (and they can be multiplied) show, it can be found also in other theoretical paradigms and other linguistic disciplines. In combination with the fact that the assumed infinite nature of language is conceptual rather than factual this shows that the role of idealisations as it was identified in the paper is not restricted to just one theoretical paradigm.

That we are dealing here with a conceptual, rather than a factual, assumption is further supported by the observation that it leads to the introduction of yet another conceptual distinction, viz., that between competence and performance. Here is an example of a standard argument that there is a specific linguistic ability that is connected with the assumed infinite nature of language (Smith, 2005, p. 31):

The important point is that humans can produce and understand any of an infinite number of largely novel sentences, and can make systematic judgments about their well-formedness, and we need to explain these abilities.

Just as actually produced linguistic material (utterances, sentences, . . .) falls short of being infinite in nature, the actual performance of a language user, too, has nothing infinite about it. Hence a separate ability, performance, must be postulated.

Again, it is important to note that although the competence–performance distinction as such is perhaps mainly known from Chomsky's work in the generative paradigm, it is an assumption that neither originates in his work nor one that is confined to that specific theoretical outlook. As Blutner notes in his comments: 'Optimality Theory [. . .] crucially relies on the competence–performance distinction'. Haspelmath makes a similar claim with regard to cognitive linguistics, stating: '[. . .] this is what most cognitive linguists do: they try

to describe the cognitive system for language (i.e., the speaker–hearers’ competence’. The distinction has roots in earlier structuralism, and in the form of an account of the ‘creativity’ of language and language users it is a well-known issue in semantics. The emphasis in semantics on compositionality reflects a similar concern.

There are several concomitant characteristics that indicate a reliance on some form of the competence–performance distinction. The assumption that the abilities that we are concerned with are individual, rather than social, is a prominent one. The view that the individual can be reduced to its mental make-up, or, as is more common nowadays, to its neural substrate, is another. Taking linguistics ultimately to be about the neural substrate of competence does not turn it into something that takes the embodied nature of language users seriously, by the way. It is rather an assertion of physicalistic monism. But even where the physicalism is lacking, the restriction to competence tends to make one’s perspective individualistic.

So even if Haspelmath is right in claiming that ‘Chomsky simply happens to lack an interest in the social function of language’, this does not mean that the focus on competence is not part of an explanation of why an entire range of theoretical frameworks has either neglected that social dimension or, at best, has tried to account for it in strictly individualistic terms. E.g., there is a principled difference between viewing communication as the sequential application of individual competences, and regarding it as a social process that essentially depends on interactions between individuals and their social and physical surroundings. The first view is internalistic in that it assumes that everything that is relevant for successful communication is contained in the individual competence, and hence, that successful communication requires complete understanding, i.e., a perfect match in the application of the competences of the communicating individuals. On the second, externalistic view, communication and complete understanding can be separated, which seems a much more plausible assumption.

2. Idealisation in linguistics

The distinction that the paper suggests between idealisation and abstraction as two fundamentally different ways of constructing suitable objects of

investigation has elicited diverse reactions from the commentators. Some regard it as a real and fruitful distinction, but argue that it is not characteristic of linguistics (e.g., Blutner), others find less merit in the distinction as such (e.g., Stabler). In this and the following sections we will discuss a couple of issues. First we address the question whether the kind of construction that we have labelled ‘idealisation’ occurs in linguistics, and follow that up by some thoughts on whether its use can be justified. Then we will go into the distinction as such and address the objection that it occurs across the board, i.e., that it is not characteristic for linguistics. We end with an attempt to clarify our use of the label ‘ideological’ in our characterisation of certain idealisations.

Let us first remark that we are not alone in drawing a distinction between abstraction and idealisation, in relation to linguistics. Here is Jackendoff in *Foundations of Language* (Jackendoff, 2002, p. 33):

Chomsky views competence as an idealisation abstracted away from the full range of linguistic behaviour. As such, it deserves as much consideration as any idealisation in science: if it yields interesting generalisations it is worthwhile. Still, one can make a distinction between ‘soft’ and ‘hard’ idealisations. A ‘soft’ idealisation is acknowledged to be a matter of convenience, and one hopes eventually to find a natural way to re-integrate excluded factors. A standard example is the fiction of a frictionless plane in physics, which yields important generalisations about forces and energy. But one aspires eventually to go beyond the idealisation and integrate friction into the picture. By contrast, a ‘hard’ idealisation denies the need to go beyond itself; in the end it cuts itself off from the possibility of integration into a larger context.

It is my unfortunate impression that, over the years, Chomsky’s articulation of the competence/performance distinction has moved from relatively soft [...] to considerably harder.

A ‘soft’ idealisation is what we call ‘abstraction’. As for the ‘hard’ variety, Jackendoff puts it rather well when he says that it ‘denies the need to go beyond itself’.

Krifka mounts a defence of idealisation, suggesting that it is inherent in Marr’s methodological proposals for cognitive science. The classical discussion of this issue is David Marr’s (Marr, 1982), where he points out that cognitive science should distinguish at least the following three levels of inquiry:

1. identification of the information processing task as an input–output function⁶
2. definition of an algorithm which computes that function
3. neural implementation of the algorithm defined

The first level incorporates a specification of the input and output representation languages. For example, in Marr's case the input representation is a two-dimensional array of (light-)intensity values, the output a mental representation of a three-dimensional shape. To be more specific, the input needs to be described as a differentiable manifold, and the output could be given in terms of a combinatorial system generating complex shapes from simple ones. In linguistics, one could specify language comprehension at the first level as: construction of a discourse model (output) from a given piece of discourse (input), and to proceed further one would have to specify also the representational formats of input and output.

Even if one doesn't know what a differentiable manifold is, the very terminology suggests that this is something one doesn't encounter in nature; but is it an idealisation (in our sense), as Krifka maintains? The underlying assumption, that intensity values are given by real numbers, i.e. with infinite precision, is clearly an abstraction, introduced to be able to represent edge detection as differentiation, but there is nothing ontological about this assumption: it is not claimed that the visual system can deal with arbitrarily precise input values. In that sense it differs from Chomskyan competence, where infinite productive capacity is a built-in feature.

Furthermore, there is no real correspondence between Chomskyan competence and Marr's informational level. The latter is the specification of an algorithm, not the algorithm itself. But competence becomes generative only once an algorithm for generation has been given, and hence competence corresponds to a combination of the first two levels. This leads to the following four observations.

First of all, it is entirely possible to have a non-productive informational level description of language, as happens for instance in constraint based grammars, where a well-formed sentence is viewed as a model of a particular first order theory. In this approach there is no need to postulate infinite productivity

⁶ Marr called this the 'computational level', but since the actual computations belong to the next level, we use 'informational' instead of 'computational'.

as the hallmark of grammatical competence. Neural networks can transform a semantic message into an utterance, which can then be checked for satisfaction of the constraints imposed by the grammar.⁷

Secondly, adopting Marr's picture leads to a softening of the hard distinction between competence and performance, since the levels mutually constrain each other; true, the algorithm is specified by the informational level, but is also constrained by the possibilities for neural implementation. Competence and performance meet in the algorithm, and it is hard to say where one ends and the other begins.

A third observation is that it is not true that taking cognitive limitations into account (e.g., those responsible for restricting the depth of embedded clauses) always leads to unwieldy algorithms. The classic counterexample is the programming language Prolog, which uses only formulae of the form $\varphi \rightarrow p$, where p is atomic and φ does not contain \rightarrow , and is therefore vastly more efficient than full predicate logic, where nested occurrences of \rightarrow are allowed.

Finally, Krifka gives a very interesting discussion of the cognitive aspects of number. He remarks, correctly, that formal systems such as Peano Arithmetic are an idealisation relative to the cognitive phenomena, and asks: 'Would Stokhof and van Lambalgen argue against [formal] arithmetic because it is not an abstraction from our cognitive abilities to deal with numbers, but an idealisation?' Krifka's own answer is that formal arithmetic allows one to make distinctions that would otherwise be inexplicable: 'We would not understand that the symbolic representation of numbers does influence the way we deal with them – for example, why the addition task $24+5$ is harder than the addition task $26+5$ in a decimal (but not in a vigesimal) system of representation. For all this, formal arithmetics is crucial'. The point here is not about arithmetic *per se*, but about the analogy between formal arithmetic and generative grammar. Our answer is that formal arithmetic certainly has a role to play in a theory of counting and its symbolic representation, but that it cannot be used to *define* counting competence. For example, this competence involves choosing ever larger 'units' to make big numbers manageable. Now, although formal

⁷ We mentioned in passing that neural syntax (in the version of Fitz (2009)) might be a useful avenue to explore, because of its ability to model recursion up to a given depth. In their comments, Gärtner and Jurish reply with a detailed critique of Fitz's work which we cannot address here, because space does not allow us to explain the formidable technicalities underlying the neural network modelling.

arithmetic has much to say about these units once chosen, in itself it does not dictate that units be chosen, since numbers are generated uniformly through iterated application of the successor function.

3. Abstraction and idealisation revisited

Both our choice of terminology and the comparison with physics have caused confusion, and have laid us open to the charge that we have an overly simplified view of modern scientific practice. In particular several commentators (e.g., Blutner, Stabler) have argued that the kind of idealisation we found objectionable in linguistics occurs in physics as well. (And hence is not objectionable? More on that later.)

First a point of terminology. The distinction between what we call idealisation and abstraction is usually made in terms of different kinds of idealisation. Jackendoff described the distinction as between ‘hard’ and ‘soft’ idealisation. In the following extended quote from Cartwright’s *How the Laws of Physics Lie* (Cartwright, 1983, p. 111) the ‘hard’ idealisations are evocatively described as ‘fictions’. The first paragraph sketches the type of idealisation that we have called ‘abstraction’.⁸ Here, some features of the phenomenon of interest are temporarily set aside, but the aim is always to re-introduce them later:

To call a model an idealisation is to suggest that the model is a simplification of what occurs in reality, usually a simplification which omits some relevant features, such as the extended mass of the planets or, in the example of the circuit model, the resistance in the bypass capacitor. Sometimes the omitted factors make only an insignificant contribution to the effect under study. But that does not seem to be essential to idealisations, especially to the idealisations that in the end are applied by engineers to study real things. In calling something an idealisation it seems not so important that the contributions from omitted factors be small, but that they be ones for which we know how to correct. If the idealisation is to be of use, when the time comes to apply it to a real system we had better know how to add back the contributions of the factors that have been left out. In that case the use of idealisations does not seem to counter realism: either the omitted factors do not matter much, or in principle we know how to treat them.

⁸ The term ‘abstraction’ is also used, but in a different sense, as that which gives physical laws generality, as when we disregard the colour of a moving object of which we compute the trajectory.

As can easily be checked this squares with the characteristic features we ascribed in the paper to abstractions. In particular the basic ‘back-and-forth’ between the case with abstraction and that without abstraction is taken to be a characteristic feature by Cartwright as well.

However, in the second paragraph Cartwright argues that this cannot be all there is to idealisation [abstraction] in physics, that there is another form of idealisation, which she calls ‘fiction’, that is of a different nature. On the face of it there is a striking analogy between such fictions in physics and idealisations in linguistics as we discussed them. But there is a crucial difference, which we will explain after the following exposition of the situation in physics:

A model is a work of fiction. Some properties ascribed to objects in the model will be genuine properties of the objects modelled, but others will be merely properties of convenience. The term ‘properties of convenience’ was suggested by H. P. Grice, and it is apt. Some of the properties and relations in a model will be real properties, in the sense that other objects in other situations might genuinely have them. But they are introduced into this model as a convenience, to bring the objects modelled into the range of the mathematical theory.

Not all properties of convenience will be real ones. There are the obvious idealisations of physics [what we have called abstractions; MS&MvL] – infinite potentials, zero time correlations, perfectly rigid rods, and frictionless planes. But it would be a mistake to think entirely in terms of idealisations – of properties which we conceive as limiting cases, to which we can approach closer and closer in reality. *For some properties are not even approached in reality. They are pure fictions.*

I would want to argue that the probability distributions of classical statistical mechanics are an example. This is a very serious claim, and I only sketch my view here as an illustration. *The distributions are essential to the theory – they are what the equations of the theory govern* – and the theory itself is extremely powerful, for example in the detailed treatment of fluid flow. Moreover, in some simple special cases the idea of the probability distribution can be operationalised; and the tests support the distributions ascribed by the theory.

Nevertheless, I do not think these distributions are real. Statistical mechanics works in a massive number of highly differentiated and highly complex situations. In the vast majority of these it is incredible to think that there is a true probability distribution for that situation [. . .]. It is better, I think, to see these distributions as fictions, fictions that have a powerful organising role in any case and that will not mislead us too much even should we take them to be real in the simple cases. (Cartwright, 1983, p. 153f.) [emphasis added; MS&MvL]

As we already noted, on the face of it there is a striking resemblance between Cartwright's fictions and our idealisations. However, there is a crucial difference between the two, which is concerned with the role that fictions and idealisation play within a theory.

By way of illustration, compare the fiction of a probability distribution over all internal states of a gas with the idealisation of infinite productivity in generative grammar, or the use of model theory in formal semantics. What differentiates between the two is that the probability distribution is said to play an organising role, but that no reality is claimed for it. The claims for infinite productivity, however, are much stronger (Smith, 2005, p. 31):

The important point is that humans can produce and understand any of an infinite number of largely novel sentences, and can make systematic judgments about their well-formedness, *and we need to explain these abilities*. The explanation in terms of merge and move may be wrong, but it is currently both the most inclusive and the deepest in terms of the deductive structure of the theory involved. [emphasis added; MS&MvL]

So whereas in physics the most profound form of idealisation – the one which involves infinities – is a useful fiction, one that helps to explain the data but that does not enter into a characterisation of the data, in linguistics the idealisation to the infinite is among the very data to be explained. This is why we believe that physics and linguistics are in the end incomparable.

To be sure, we do not argue that in linguistics no abstractions should be used, as, e.g., Haspelmath seems to think when he writes: 'While they do not say so explicitly, it seems that the conclusion is that students of language should not use abstractions at all, but should study their observables in all their empirical messiness'. This is not our position. Abstraction from 'empirical messiness' is needed in any systematic study, and linguistics is no exception to that rule. Our point is, rather, that because language is a phenomenon that has such a diverse ontological status, – part physical/biological, part social/cultural/historical – , one should be careful in what one abstracts away from, so as to not turn what might look like an abstraction into what actually is an idealisation.

This is not to suggest that all idealisations are merely botched attempts at abstraction. In some cases idealisations are the result of the attempt to adhere to a particular methodology for reasons that are not intrinsically connected with the subject matter at hand. In such cases idealisation is the result of what we have called 'ideological motives'.

4. Idealisation and ideology

So, finally, a word about what we have called the ‘ideological’ nature of idealisation, a term that has led to some confusion. What we mean when we call some idealisations ‘ideologically’ motivated is that the background motivation for making them in the first place seems to stem, not from observations and concerns regarding the empirical phenomenon that is at stake, but rather from the adoption of a particular philosophical point of view on, for example, the nature of scientific investigation, or the nature of reality, or that of knowledge, or, as is usually the case, some combination thereof.

A nice illustration is provided by Stabler’s reference to Plato’s notion of ‘carving nature at its joints’. The idea of ‘Nature’ as a given, with natural, independent distinctions between its various parts, is as such a good example of a philosophically motivated, and hence, in the present terminology, ‘ideological’ assumption. That does not mean that it cannot be a useful assumption to make. The point to bear in mind is that it is not the philosophical motivation *per se* that makes such an assumption ideological, it is rather a lack of empirical justification for a particular application that is responsible for this.

Needless to say, ideologically motivated assumptions and concepts are widespread, and are certainly not characteristic of the kind of concerns in linguistics that we address in the paper. Evidently, the scare quotes we used in the paper when discussing differences between linguistics and ‘natural sciences’ were not enough to prevent such a misinterpretation. So, once more: we do not believe that in the sciences idealisations do not occur. The work of Cartwright and of many others in philosophy of science and in science studies has shown, ideology also informs the sciences, at some points, and to a certain extent. However, that does not mean that the use of such motives is innocuous, nor that its consequences should not be investigated, also in linguistics.

In the paper we focussed on what we there called the ‘naturalistic presupposition’ that is characteristic of the generative tradition. In hindsight, our use of the term ‘naturalistic’ was perhaps less fortunate and somewhat confusing. It is not that there aren’t elements of naturalism in the generativist framework, but they do tend to be interpreted in quite a specific manner, one that is informed by a particular presupposition of what a proper scientific investigation of any empirical phenomenon should look like.

Intuitively, a naturalistic approach to natural language would be one which takes it as a natural phenomenon that has both biological/ psychological and

social/historical dimensions. That, however, is not the kind of naturalism that generativism, and in its wake many other paradigms in modern linguistics, subscribe to. Rather, it is a decidedly physicalistic kind of naturalism, one that seems to be committed to the implicit assumption that, ultimately, a theory of human language that is descriptively and explanatorily satisfactory, needs to appeal to nothing beyond the realm of the physical make-up of humans. This kind of naturalism is a-historical, individualistic and thoroughly reductionistic in its outlook on reality. Also, it takes the deterministic, covering law model of explanation that many still hold as characteristic for the sciences as the sole model of true scientific explanation, and it will regard anything that does not live up to these standards as inherently flawed.

5. The actuality of the analysis

There is one issue we haven't touched upon in the above, and that is Newmeyer's argument that we have misrepresented the status of theoretical linguistics as a role model for the humanities. According to Newmeyer the dominance of various forms of postmodernism is inversely proportional to the intellectual prestige of theoretical linguistics, in particular of the generative kind. Also with regard to other fields, in particular cognitive science, Newmeyer discerns a steady decline of the influence of the generative paradigm. Two comments are in order here. First of all, our statements that Newmeyer disagrees with referred to the role of linguistics in the early days, i.e., the sixties and seventies of the last century, i.e., before various forms of post-modernism became fashionable in the humanities, and when the 'cognitive revolution' was still in its infancy. After that, indeed, a lot has happened, and the landscape has changed considerably.

Whether the present state of generative linguistics is indeed as sorry as Newmeyer describes it, depends perhaps on what parameters one looks at: numbers of scholars engaged, *de facto* relevance for other fields, reputation among non-scholars. Our own experience, for what it's worth, is that among scientists and intellectuals in general, theoretical linguistics still enjoys great prestige, precisely because in other disciplines in the humanities various post-modernistic methodologies are still, if not dominant then at least most visible. Linguistics enjoys prestige because it appears to be conducted in a manner that most people associate with what they regard as the only proper one, viz., that

of the sciences. This judgement may be unjustified, but it is a matter of fact that it still exists. Whether it will prevail is, of course, anybody's guess. But the continuing dominance of a scientific mode of thought, especially in those institutions that ultimately determine the financial future of scientific research, should give pause for thought.

*Correspondence address: ILLC/Department of Philosophy
Universiteit van Amsterdam
m.j.b.stokhof@uva.nl
m.vanlambalgen@uva.nl*

References

- Cartwright, Nancy. 1983. *How the laws of physics lie*. Oxford: Clarendon Press.
- Fitz, Hartmut. 2009. *Neural syntax*. Amsterdam: ILLC/Department of Philosophy dissertation.
- Hauser, Marc D., Noam Chomsky & Tecumseh W. Fitch. 2002. The faculty of language. *Science* 298. 1569–1579.
- Jackendoff, Ray. 2002. *Foundations of language*. Oxford: O.U.P.
- Karlsson, Fred. 2002. Is there an upper limit to right-branching embedding of clauses. In *Publications of the Department of General Linguistics*, 196–199. Tartu: University of Tartu.
- Lasnik, Howard. 2005. Grammar, levels, and biology. In James McGilvray (ed.), *The Cambridge companion to Chomsky*, 60–83. Cambridge: C.U.P.
- Marr, David. 1982. *Vision*. New York: Freeman Press.
- Montague, Richard. 1970. Universal grammar. *Theoria* 36. 373–398.
- Pinker, Steven & Ray Jackendoff. 2005. The faculty of language: What's special about it? *Cognition* 95. 201–236.
- Pullum, Geoffrey & Barbara Scholz. 2010. Recursion and the infinitude claim. In Harry van der Hulst (ed.), *Recursion in human language*, 113–138. Berlin: Mouton de Gruyter.
- Rutherford, Melissa D. 2004. The effect of social role on the theory of mind reasoning. *British Journal of Psychology* 95. 91–103.
- Smith, Neil. 2005. Chomsky's science of language. In James McGilvray (ed.), *The Cambridge companion to Chomsky*, 21–41. Cambridge: C.U.P.
- Tesar, Bruce, Jane Grimshaw & Alan Prince. 1999. Linguistic and cognitive explanation in Optimality Theory. In Ernest Lepore & Zenon Pylyshyn (eds.), *What is cognitive science?*, 295–326. Oxford: Blackwell.
- Tomalin, Marcus. 2006. *Linguistics and the formal sciences*. Cambridge: C.U.P.