

Abstractions and idealisations: The construction of modern linguistics

MARTIN STOKHOF and MICHIEL VAN LAMBALGEN

Abstract

The paper addresses the way in which modern linguistics, – in particular, but not exclusively, the generative tradition –, has constructed its core concepts. It argues that a particular form of construction, reminiscent of, but crucially different from, abstraction, which is dubbed ‘idealisation’, plays a central role here. The resemblances and differences between abstractions and idealisations are investigated, and consequences of the reliance on idealisations are reviewed.

1. Introduction

In many ways, modern linguistics is one of the most remarkable and successful scientific innovations of the twentieth century. The rise of generative grammar in the fifties and sixties produced an atmosphere of intellectual excitement that seemed to be reserved for fundamental developments in the natural sciences. And the excitement was not restricted to linguistics as such, it stretched out to other disciplines, such as philosophy, the emerging disciplines of computer science and cognitive psychology, anthropology and literary studies. And to the present day modern linguistics is held up as a model of scientific innovation to other disciplines in the humanities.

A satisfactory account of this remarkable development will have to factor in a number of things. The role of the natural sciences and the formal sciences as a ‘standard model’ of scientific inquiry is one of them. Another is the way in which modern linguistics appears to tie in with internal, disciplinary developments in other fields. Sociological factors, such as the way in which the

discipline organises itself, are also relevant.¹ And then there is the way in which linguistics appears to have succeeded to conceptualise its central objects of study so as to fit a particular methodology.

In this paper we deal with this last issue, i.e., with the question how modern linguistics has constructed its objects of study, such as ‘language’, ‘grammar’, ‘competence’, ‘meaning’, ‘rule’. Apparently, a major factor that explains the success and prestige of modern linguistics is that it has succeeded to come up with scientific characterisations of its core concepts that have allowed linguists to develop theories that are both descriptively and explanatorily adequate. In what follows we focus on a particular aspect of this complicated process that, we feel, has not received adequate attention in the literature to date, viz., the nature of the kind of constructions that modern linguistics employs.

There are two things we would like to mention at the outset. First of all, in what follows we use the phrase ‘modern linguistics’ mainly as an indication of what is still a dominant approach, viz., the generative tradition. And secondly, our considerations primarily have a ‘meta’-character, i.e., the observations that follow are not intended as arguments pro or con particular positions, although they could have such repercussions. But the spelling out of such consequences is beyond the scope of this paper.

2. The state of the art

As we noted above, the rise of modern linguistics, its success and influence, and its enormous intellectual prestige, as such are intriguing phenomena that call for an explanation. But also from an internal perspective, i.e., from the perspective of linguistics itself, its present state is one that raises a number of questions.

One of these is, that despite the solid reputation that linguistics has as a successful discipline, many of the expectations have not (or not yet?) been realised. If we look at the description of individual languages, we can note that complete and explicit grammars are still far off. In the area of typology

¹ A thorough, empirical sociological study of the development of modern linguistics does not exist, as far as we know. For studies that are more of the nature of a ‘history of ideas’, cf., e.g., Newmeyer (1986), Harris (1993).

many studies have been done, but it remains to be seen how much of that work actually depends on the methodology of modern linguistics. Little or no explanations of properties of natural languages exist that are accepted generally, i.e., across theoretical boundaries. When it comes to applications, especially computational ones, we can observe that the theoretical models of modern linguistics, based as they are on the concept of a grammar as a rule system, in general are less successful than stochastic approaches. And with regard to psycholinguistic investigations and research into the neurophysiological processes that underlie language and language use, it appears that modern linguistics in general is unable to come up with leading questions and hypotheses.²

Another observation regarding the present state of modern linguistics, and one that definitely calls for further study, is the substantial diversity in approaches and models, and even in definitions of central concepts, that has become a distinctive feature of linguistics to date. With the rise of generative grammar, as proposed and developed by Chomsky and others, modern linguistics seemed to be heading towards a remarkable uniformity vis à vis its goals, methodology, and central concepts. At least this appeared to hold for core disciplines such as syntax, morphology, and phonology. In semantics a similar development occurred at the end of the sixties when formal semantics appeared on the scene. ‘Montague grammar’ apparently developed into a generally accepted model for semantic description and explanation. But the uniformity and consensus that at some point seemed almost natural have disappeared: there is an enormous variety of approaches, theoretical models, methodologies, and even with regard to the goals of linguistics and its very object of investigation there are fundamental differences of opinion.³

These observations give rise to a fundamental question with regard to linguistics as such: Could modern linguistics perhaps be an example of a ‘failed discipline’? As was already noticed above, the adoption of the models and methodologies of the natural sciences and the formal sciences was one of the keys to the success of modern linguistics. Moreover, especially in Chomsky’s

² To be sure, this is not just a problem for modern linguistics. Quite generally, it is difficult to derive from theories concerning macroscopic phenomena predictions regarding the underlying neurophysiology due to the absence of clear bridging principles that link the often disjoint conceptual systems.

³ Cf., Kamp & Stokhof (2008) for a description of this development, and an attempt to explain what drove it, for the case of formal semantics.

views a clearly naturalistic goal can be discerned: according to him linguistics studies what in the end is an aspect of human biology. Is this naturalism perhaps one of the causes of the present, confusing situation? Is it that modern linguistics, knowingly or unknowingly, follows a naturalistic approach to phenomena – language and linguistic competence – that are of a fundamentally different nature?

This last question is too complex to be even properly articulated in the context of this paper, let alone that it can be answered here. However, we do feel that the observations about abstraction and idealisation as constructive processes that are the subject of what follows do present reasons to think that the question just formulated touches on a central problem with regard to the status of modern linguistics as a scientific discipline. And if we are correct in thinking so, then it is also the case that, precisely because modern linguistics has functioned as a model for other disciplines in the humanities for more than four decades, the relevance of this question extends beyond linguistics as such.

3. Examples of constructions

To give the reader some idea of the kind of constructions⁴ we have in mind, here are a few examples.

At first sight, ‘language’ appears to be the most central concept of linguistics. Be it specific natural languages, such as English or Quechua or Renellese, or natural (human) language in general, language seems to be the core phenomenon that linguists want to describe and explain. Now, from an observational point of view language is first and foremost language use: spoken or written utterances.⁵ For the child that acquires its mother tongue, language use is what it encounters in its environment, for adult language users language is what they use to communicate with each other.

⁴ What follows will make clear that the term ‘construction’ is used here not in its linguistic sense, but as a term that belongs to the vocabulary of philosophy of science.

⁵ Obviously, spoken language is primary vis à vis written language, not just historically but also ontogenetically. Yet in linguistics, as is the case in most philosophical treatises on language, the focus is mainly on written language, not on speech. Cf., Kraak (2008) for a recent study of the effects of this shift.

In modern linguistics⁶ the intuitive concept of language, viz., that encountered in everyday use, has been replaced by the logical, mathematical (algebraic) concept of a language, viz., that of a potentially infinite set of well-formed expressions generated by a finite, or finitely characterisable, set of rules (i.e., a grammar).⁷ Not only does this concept emphasise the formal aspect of language, and hence the focus on written language, it also introduces a notion of ‘structure’ that can be tested against actual linguistic material only indirectly, and partially.⁸ Another immediate consequence of the shift towards a formal construction of the concept of language is that expressions are being studied at the level of types, not tokens, with regard to both their form as well as their meaning. Obviously, the historically contingent availability of writing is instrumental in this change.

A related move is that linguistic competence, i.e., the ability of humans to use language, actively in production and passively in interpretation and understanding, is being studied in terms of a comparable construction. Here the well-known distinction between ‘competence’ and ‘performance’ plays a key role. Knowledge of a language is conceived as the availability of a grammar, and competence as the ability to use that grammar to distinguish well-formed expressions from non well-formed ones, to assign the former an interpretation, and then to use them both actively and passively. This linguistic competence, though an individual capacity in the sense of being ascribable to an individual as such, is not introspectively accessible to the individual that has it.⁹

⁶ What follows applies not just to the generative tradition, but also the many approaches it has helped shape in this respect. But there are other approaches in which the construction described here does not play a role, or at least not in the same way.

⁷ Cf., Tomalin (2006) for an extensive study of the role that the developments in logic in the first half of the twentieth century have played in Chomsky’s early work.

⁸ In the light of this, one particular development in modern linguistics becomes more easy to understand, viz., the fact that one of the most central notions, that of ‘syntactic structure’, has been subjected to many, and radical, changes. This constant re-conceptualisation and re-modelling of a core notion makes sense only if we keep in mind its mainly theoretical nature (and that of related notions, such as ‘rule’, ‘constituent’, and so on). Cf., Stokhof (2002) for a discussion of similar observations with regard to the central notion of semantics, viz., ‘meaning’.

⁹ This creates what Jackendoff (1987, p. 20) calls ‘the mind – mind problem’. On the one hand, we can be clearly and consciously aware of what we do with language (we may consciously opt for a certain interpretation of an utterance, or for a certain formulation of what we want to say, we may be at a loss as to the meaning of what is being said, or object to a certain

Another phenomenon, that is closely related to the idea of competence as an individual ability and that has strongly influenced contemporary thought about language, and hence also the goals of modern linguistics, is the so-called ‘problem of creativity’ (or ‘compositionality’). It is the ‘observation’¹⁰ that a language consists of a potentially infinite number of well-formed expressions that somehow has to be represented in a finite manner in the finite individual human brain. In a certain sense this ‘problem’ is generated directly by the shift towards the logical, mathematical characterisation of the core concept of a language. Closely related is what Kraak in his aforementioned book calls ‘the myth of representation’, viz., the idea that language, and in particular written language, serves as a medium of representation of internal, mental contents. If we assume that humans are capable of a potentially infinite number of thoughts (and desires, and conjectures, and questions, and so on), then the myth of representation inevitably leads to the conclusion that the language we use to express such contents also has to have an unlimited character.

These constructions, and others like them, lead to a relative neglect of both the actual use of language as well as the context in which that actual use appears: the physical, social and cultural environment, both synchronically as well as diachronically. Whenever attention is being paid to language use, it is always as complementary to the idea of language as characterised by the form and (literal) meaning of its expressions. Almost all theories about what it is that people do with language start from these very assumptions about what language and linguistic competence are. The result is very much an abstract and individualistic picture: linguistic competence is an individual ability, and language use is a process in which autonomous and competent individuals exercise their linguistic competence. That language use has a social nature, in which communication plays a central role, is, of course, not something that many linguists would like to deny. But, so the leading idea proclaims, the language that is being used and the competence that is being applied in that social process, can be described, characterised and explained as such, and quite inde-

choice of words for a number of reasons), but, on the other hand, the mechanisms that are postulated to constitute the essence of our competence are in principle shielded from direct inspection.

¹⁰ In scare quotes because in fact of any natural language only a finite number of utterances will ever be observed. Cf., Groenendijk & Stokhof (2005) for further discussion.

pendently from language use.¹¹ Behind this is the fundamental assumption that in the end language and linguistic competence can be understood as phenomena that are anchored in human biology, and that it is only via the methodology of the natural sciences that we may acquire insight into their nature and function.

This, admittedly concise, sketch of some core moves in the construction of the central concepts and goals of linguistics gives reason to believe that modern linguistics has been decisively influenced by ideas and developments in other disciplines, notably the formal and the natural sciences, but also philosophy. As for the influence of the latter, Chomsky's rationalism is an obvious and explicit example, but at other points it is more subtle and therefore perhaps less often noticed.¹² In what follows we will not so much be concerned with the actual details of such constructions, but rather focus on the nature of the process as such. In doing so, our central question is the following: Are these constructions like the abstractions we are familiar with from the natural sciences, or are they of a different nature? And if the latter turns out to be the case, what are the consequences for the status of linguistics?

4. Abstractions as constructions

Abstraction is a well-known tool for turning a natural phenomenon into a 'suitable' object of scientific investigation. Standard examples are the frictionless plane in classical mechanics, the perfect vacuum, pure chemical substances, and so on. Whereas in reality moving objects always are subject to friction, a perfect vacuum does not exist and cannot be created, and chemical substances almost always contain contaminations from other substances, these facts, when considered from the point of view of studying certain central natural phenomena, are complications which are either deemed irrelevant or too complex or intractable to be captured in a theory, at least for the time being. The latter phenomena in particular are interesting if we want to determine what exactly it is that an abstraction is, and does.

¹¹ This is very much the dominant view, one that can be found explicitly in the work of Chomsky, and one that has gone unchallenged for a long time. Recently other views have started to emerge. In the concluding section we will briefly mention some of them.

¹² Cf., Stokhof (2002, 2007, 2008) for an analysis of various philosophical distinctions and goals that have shaped and continue to guide formal semantics.

The physical theory of tides provides another illustrative example. Newton's theory of 1687 gave an explanation of the frequency and amplitude of tidal waves based on his theory of gravitation, in terms of the combined gravitational pull on the earth exercised by the sun and the moon. His calculations assumed that the entire surface of the earth is covered by one ocean and that this ocean has no inertia of its own. These two assumptions meant that, first of all, local circumstances on the earth could not play a role, and, second, that the earth's rotation was not taken into account. Also, the effect of other celestial bodies, such as the planet Venus, was disregarded.

Of course the reality of the phenomena that did not fit into this model was not denied. In fact, further work on the theory produced a model in which these phenomena can be accounted for, using both physical calculations as well as observations of the local circumstances at locations where the actual tidal heights needed to be calculated. (Relevant factors include the depth of the ocean, the form of coast lines, the presence of pack ice, and so on.) The more accurate model is analogous to that of a vibrating violin string: the timbre of the sound it produces is determined by the many frequencies, each with its own amplitude, that co-occur with the basic tone. Analogously, the periodic process of tidal waves is determined by many frequencies, some of which are determined by astronomic laws, others by local circumstances.¹³ But even in this more complex model one is forced to abstract, since some frequencies, such as the disturbances caused by moving sand banks, are too difficult to predict. However, the reality of the factors from which one abstracts, is never denied, and in principle the model is capable of incorporating them.

This is a crucial feature of the way in which abstraction in the natural sciences works: the phenomenon from which we abstract is a real one, and its reality is acknowledged in the theory or in the model that is based on the abstraction. After all, in factual observations and experiments these phenomena inevitably occur. One of the main reasons for nevertheless abstracting from them is that by doing so one is able to come up with a better explanation of the underlying causal mechanisms while keeping the predictions of the theory based on the abstraction within certain acceptable limits of accuracy.

¹³ In fact, in the case of tidal waves the model is somewhat more complex, since, unlike in the case of the vibrating string, there is more than one basic tone at work at the same time.

This means that there is a real and acknowledged interaction between the theory, i.e., the explanation it provides of a certain phenomenon together with the predictions it delivers, and reality as it occurs in observations and experiments. Another example of this is provided by the concept of a perfect vacuum. In physics so-called ‘free space constants’, such as the speed of light and the magnetic constant, play a key role. The quantitative values of these constants is theoretically determined with reference to a perfect vacuum. In reality, in which a perfect vacuum does not occur, these constants always have slightly different values, but the differences can be approximated with sufficient precision to make the predictions of the theory practically useful. (And in many cases the differences are so small that they can be safely ignored.) So what we see is that theory based on abstraction and observation and experiment without abstraction remain intimately connected, both conceptually as well as practically.

And the reason that this is a crucial feature of the way in which abstraction in the natural sciences works is that it explains why theories that make use of abstractions still work: they do not ‘re-conceptualise’ the phenomena.

5. Abstraction in linguistics?

In modern linguistics, too, we often find appeals to abstraction when it comes to explaining how a linguistic theory is related to observable reality. The following quote from Chomsky (1980, p. 219) illustrates what is at stake:

Any serious study will [. . .] abstract away from variation tentatively regarded as insignificant and from external interference dismissed as irrelevant at a given stage of inquiry. [. . .] It should come as no surprise, then, that a significant notion of ‘language’ as an object of rational inquiry can be developed only on the basis of rather far-reaching abstraction.

What Chomsky is suggesting here is that abstraction in linguistics is the same process as in the natural sciences. It allows us, he claims, to concentrate on the core of the phenomenon, disregarding those aspects that are deemed ‘insignificant’ or ‘irrelevant’. As such this is a remarkable statement, because as we have seen above, in the natural sciences abstraction usually does not concern irrelevant or unimportant aspects of phenomena, but features that for one reason or another cannot (yet) be incorporated into the theory because they are

too complex or intractable. Note also that in this passage Chomsky does not provide any argument why for example the phenomenon of language can be studied only via abstraction. What is it that he means by a ‘serious study’ or a ‘rational inquiry’ that it can only be done on the basis of far-reaching abstractions?¹⁴

For Chomsky, then, it is apparently obvious that the fact that language and linguistic competence, certainly at first sight, are different kinds of phenomena than movement of physical bodies or chemical reactions, constitutes no reason to think that abstraction could not, and should not, play the same role as it does in the natural sciences. Thus he writes in Chomsky (1995, p. 7):

... it is a rare philosopher who would scoff at its [i.e., physics’] weird and counter-intuitive principles as contrary to right thinking and therefore untenable. But this standpoint is commonly regarded as inapplicable to cognitive science, linguistics in particular. Somewhere between, there is a boundary. Within that boundary, science is self-justifying; the critical analyst seeks to learn about the criteria for rationality and justification of scientific success. Beyond that boundary, everything changes; the critic applies independent criteria to sit in judgment over the theories advanced and the entities they postulate.

But this really rests on a misrepresentation of how things are done in the natural sciences. No physicist, for example, would be of the opinion that any aspect of a physical theory is ‘self-justifying’, including the abstractions on which the theory is based. The final judgement always resides with observational and experimental verification and explanatory adequacy. In other words, the last word is spoken, not by the physicist (and, of course, also not by the philosopher), but by reality itself.¹⁵

Apart from this misrepresentation, what is intriguing about this passage is that Chomsky apparently thinks that criticism of the constructions that define modern linguistics is not justified because the mechanism employed there does not differ from that in the natural sciences. To put it differently, Chomsky does

¹⁴ For an incisive criticism of Chomsky’s often heavily rhetorical writing, cf., Paul Postal’s essay ‘Junk Ethics’ in Postal (2004, Part 2).

¹⁵ No doubt there are concrete instances in the development of the natural sciences where one might observe a difference between ideology and practice, e.g., when empirical observations are neglected in favour of a theoretically motivated judgement. But that is not what is at stake here. What counts is that in the end one is willing to let the facts, such as they are to the best of one’s knowledge, have the final say. And that principle stands also in the case of theories that are founded on abstractions.

not differentiate criticism of the process from criticism of the result. But the question is whether that is justified in this particular case. In order to see whether it is, we take a somewhat more systematic look at the essential features of abstraction in the next section.

6. Features of abstraction

There is some discussion in the literature about the role of abstraction in the natural sciences,¹⁶ but that by and large concentrates on the modelling of this mechanism (in terms of formal models of theories, theoretical vocabularies, and so on). Though interesting and important, these are not the aspects we are concerned with here. Our primary interest concerns those features of abstraction that may settle the question whether abstraction plays, or should play, a role in linguistics.

From the examples we have briefly discussed in section 3 the following features of abstraction emerge:

- Object: a quantitative parameter of a phenomenon that is subject to abstraction, is assigned a specific value (zero, infinite, . . .)
- Result: a model of a phenomenon in which the parameter that is being abstracted over is still present
- Motivation: primarily methodological and practical

The quantitative nature of the object of abstraction does not come as a surprise: most theories in the natural sciences aim for a description and explanation of phenomena in terms of interactions and causal connections between quantitative features (speed, mass, spin, magnetic force, and so on). Relevant candidates for abstraction then are those quantitative features of which the exact actual values are irrelevant or too complex to determine. Examples of the former are the exact values of the afore-mentioned physical constants, keeping in mind that the question of '(ir)relevance' ultimately depends on the application of the theory. Examples of the latter we may find for example in the theory of tidal waves, in fluid dynamics and in the study of other semi-chaotic physical systems.

As for the result of abstraction, what is crucial there is that abstraction is not the same as negation. What is being neglected is the actual value of a parameter

¹⁶ Cf., e.g., Jones (2005).

in a concrete situation, but *not the parameter itself*. For example, if we employ the concept of a perfect vacuum we assume that there are no particles with mass, but not that mass is not a relevant concept.¹⁷ In this sense abstraction is conservative: in the resulting model the features that we abstract over are still present. In other words, abstraction does not change the ontology of the phenomena, and that makes it possible, at least in principle if not always in practice, to ‘undo’ an abstraction. This is also evident from the fact that the predictions we derive from a theory based on an abstraction can actually be compared with observations and the outcomes of experiments.

And in the end, that is what we actually want, since it is only through observation and experiment of the phenomena as they actually present themselves that we can evaluate our theories and gauge their explanatory power. In other words, abstraction first and foremost is a means to an end, it is there to enable us to start theorising by lifting some of the epistemological burden. In sum: abstraction is methodologically and practically motivated, not ontologically or ideologically.

7. Features of idealisation

As we will illustrate in this section, the type of construction that is used in linguistics and that is often taken for abstraction as it is used in the natural sciences, differs from the latter on a number of fundamental points. In particular, in linguistics the objects lack the quantitative nature that is so characteristic for objects of abstraction in the natural sciences. What we are dealing with in linguistics are rather qualitative features of phenomena that are being ignored. In order to terminologically distinguish the two types of construction we will reserve the term ‘abstraction’ for the process that we know from the natural sciences, and use the term ‘idealisation’ to refer to the kind of construction that occurs in linguistics.¹⁸

Distinctive features of what we call idealisations are the following:

¹⁷ As another example, cf., how negation functions in the law of inertia: ‘If the vector sum of all forces (that is, the net force) acting on an object is zero, then the acceleration of the object is zero and its velocity is constant’.

¹⁸ Do note that both terms, ‘abstraction’ and ‘idealisation’, are used in the literature also in other ways. Cf., the afore-mentioned Jones (2005).

- Object: a qualitative feature of a phenomenon that is being ignored
- Result: a model of a phenomenon in which the feature that is being idealised is missing
- Motivation: primarily ideological and theoretical

One of the reasons that idealisation differs from abstraction is that whereas the objects of study in the natural sciences are defined (mainly) quantitatively, those in the humanities are (primarily) characterised in qualitative terms. A definition of, say, ‘epic poetry’, or ‘the western christian tradition’, but also of such objects as ‘meaning’ or ‘subject’, determines an object (almost) completely in terms of qualitative properties. Consequently, a scientific study of such objects focusses on those properties and their relationships with other, similarly qualitative features. Quantitative features (such as determinations of time, location, and so on) may play a role also, of course, but usually they are not really essential, neither for the definition of the object of study as such, nor for the explanations that one is after. What is important to note is that leaving one or more of such qualitative features out of consideration, is not abstraction in the sense in which we discussed that in the previous section. It does not concern a quantitative parameter the value of which is fixed, but a qualitative feature that is left out.

One consequence of this fundamental difference is that the result of an idealisation is likewise fundamentally different from that of an abstraction: in the resulting model the phenomenon in question has turned into something essentially different from the original one. In other words, in the case of idealisation we are dealing with an ontological change, rather than with an epistemological one, as is the case with abstraction. Obviously, this has repercussions for the relation between the idealisation and the original phenomenon: that relation is not longer ‘symmetrical’. A simple example may serve to illustrate the point. If in a study of the western christian tradition one limits the object of study to the church, and leaves out aspects that are related to lay people, lay communities and the like, then one actually studies a different (in this case, more restricted) phenomenon, and one cannot expect that explanations and connections that are uncovered in the limited model extend to the broader phenomenon. In fact, the limited model will simply not make any predictions ‘beyond its scope’ whatsoever.

The motivation for a particular idealisation may very well be practical in nature (as in the simple example just given), and as long as one remains aware

of the implied restrictions it may be an unobjectionable move. However, quite often the motivation is not so much practical as ideological. Then certain features of a phenomenon are left out because one wants to apply a specific methodology to the idealised result. That is a move that is based on ideological reasons having to do with the conviction that only certain methods lead to scientifically reputable results. As we will illustrate below, idealisations in linguistics are often motivated by such ideological concerns.

It is worth noting that methodological considerations may play two, essentially different roles. In some cases the choice for a particular methodology is justified by an assessment that the use of a particular method increases the chance of a successful investigation, where what counts as ‘successful’ is determined *independently* of the methodology as such. But one may also choose a particular methodology on ideological grounds, in which case what counts as ‘success’ *is changed by the methodological choice* (partly because it changes the nature of the object of study). Abstraction is a methodological choice of the former kind, idealisation, in so far as it is (also) motivated by ideological concerns, one of the latter.

8. Idealisation in linguistics: an example

One of the most prominent and well-known examples of construction in modern linguistics is the ‘competence – performance’ distinction. In his groundbreaking book *Aspects of the Theory of Syntax* (Chomsky, 1965, p. 3) Chomsky introduces the distinction in the following way:

Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest and errors (random or characteristic) in applying his knowledge of the language in actual performance.

What happens here is that competence, regarded as the proper object of study of linguistics, is constructed from what we can observe, i.e., everyday use of language, by stripping it from a number of features, such as memory limitations, mistakes, (communicative) goals, attention shifts, and so on. In other words, Chomsky constructs from observable language use a concept of linguistic competence by simply ignoring a number of its actual, real properties.

In that way a new object of study is created, i.e., an object that has an ontological status that differs from that of the original one.

The reasons for this construction are not given in the passage quoted, it is just being asserted that the features that are left out by the idealisation are ‘grammatically irrelevant’. In other words, it is claimed that in the study of language, grammar, and linguistic competence, no attention needs to be paid to such factors as memory, attention, goals, and the like. But note that this claim does not rest on a comparison of (the study of) two independently given objects, viz., idealised competence and actual language use. Rather, one of the two, competence, is being constructed on the basis of this claim, and hence whatever results studying it provides cannot give independent evidence that justifies the construction in the first place. This is a strong indication that we are dealing with an ideologically motivated claim.

That this idealisation actually creates a new object is also evident from the fact that the relation between the original phenomenon, of observable language use, and the new idealised object, competence, creates *new* issues:

To study actual linguistic performance, we must consider the interaction of a variety of factors, of which the underlying competence of the speaker-hearer is only one.

This passage, also from *Aspects* (p. 4), illustrates that an idealisation raises additional epistemological questions, viz., how the idealised object and the original observable phenomenon can be related to each other. This is quite different in the case of an abstraction, where the relation between the abstraction and the phenomenon really boils down to a specification of actual values of quantitative parameters, a procedure that, though sometimes hard to carry out in practice, does not introduce any new epistemological problems.

This complication is a real one. For example, if we construct a competent language user by idealisation as an individual with implicit knowledge of the grammar of his/her *I-language*,¹⁹ we leave out many of the features that are characteristic of actual language users: the already mentioned memory limitations, the fact that language is used in order to reach certain goals (most of the time non-linguistic ones), the social environment in which language is used, but also for example the fact that language users are embodied subjects. Such factors, precisely because they are ‘idealised away’, are no longer present in

¹⁹ Cf., Chomsky (1986) for the introduction of the concept of ‘I-language’. For a thorough criticism from a broadly Wittgensteinian perspective, cf., Stein (1997, chapter 3).

the model of the competent language user (a competent user is ‘disembodied’), and *the resulting model by itself does not contain any suggestion or clue as to how it could be related in the end to what we can in fact observe*. In that sense, idealisations don’t make life any easier, on the contrary, they create a lot of extra work.²⁰

9. Abstraction versus idealisation: characteristics and backgrounds

Abstraction and idealisation, then, are two radically different ways in which objects of scientific investigation can be constructed. In Table 1 we summarise their various characteristics:

Table 1. *Abstraction versus Idealisation*

Abstraction	Idealisation
methodological	ontological
symmetric	asymmetric
no ontological consequences	additional epistemological tasks
quantitative	qualitative

An obvious question is why abstraction works in the natural sciences, but not in linguistics. It appears that this is no coincidence but something that is intimately related with the nature of the respective enterprises and with the nature of their respective domains of inquiry.

By way of illustration Table 2 lists some differences between research in the natural sciences and research in linguistics that are pertinent to this issue:

Table 2. *‘Natural science’ versus Linguistics*

‘Natural science’	Linguistics
experimental design	hardly any experiments
natural ontology	hybrid ontology
quantitative differences between theory and application	qualitative differences between theory and application
no ontological consequences	additional epistemological tasks
deterministic explanation	interpretative explanation
causal laws	no strict laws

²⁰ Another illustration of this effect, connected with the construction of an ideal competent user, has to do with the ‘knowledge’ such a user is supposed to have of language. The postulated

That the natural sciences²¹ are intrinsically based on an experimental design is closely related to the symmetric nature of the relation between a theory based on abstractions and the natural phenomenon we investigate via observation and experiment. It is due to the experimental design that there are the necessary ‘checks and balances’ on the relation between theory and practice, and due to their quantitative nature abstractions respect those constraints. In its turn, this relates to the primarily methodological nature of abstractions: they do not change the nature of the object of study. This means that a theoretical prediction can be tested by means of an application on the original, natural phenomenon, precisely because the parameter which the abstraction fixes at a certain value has been preserved in the theory. Linguistics lacks an experimental design, and hence everything that comes with it.²²

Unlike abstractions, idealisations are not methodological but ontological in nature. They change the object of study, and one of the consequences of this is that there no longer is an immediate relation between the idealised object and the original, natural phenomenon. And that means that predictions derived from the theory cannot, at least not as such, be tested by means of an application to the phenomenon. We always need an additional ‘bridging’ theory that connects the idealised object and the natural phenomenon. Not only is creating such a bridging theory an additional epistemological task, also, because of the theoretical nature of the idealised object, it is very hard to base such a bridging theory on empirical data. And that compromises the empirical nature of the theory based on the idealisation as such.

A possible, and we think plausible, explanation of this difference between natural science and linguistics comes from the nature of their respective ontologies. The natural sciences deal with ontologies consisting of natural phenomena that are subject to strictly deterministic²³ causal laws that can be

mental state lacks several characteristic features of what knowledge is, and hence requires the introduction of yet another idealised concept: the competent user ‘cognises’ language.

²¹ ‘Natural science’ here represents a number of central characteristics of various disciplines, such as physics, biology, chemistry, and so on. Of course we are well aware that a characterisation of the differences between various fields of science is an enormously complicated and, at points, questionable enterprise, and that what is listed in Table 2 needs to be extended and nuanced in many ways. However, for our present purposes this rough indication suffices.

²² To be sure, in psycho(patho)linguistics experiments are being conducted, but these are (almost) never experiments that attempt to test two alternative linguistic theories.

²³ We disregard the indeterministic nature of quantum physical phenomena, because that is not relevant for the issue at hand.

formulated in quantitative terms. Linguistics, on the other hand,²⁴ is concerned with an ontology that is not purely natural in the same sense. The phenomena that linguistics studies admittedly have physical, biological, and psychological features, but at the same time they are also historical, social and cultural phenomena. It is the hybrid nature of the ontology that explains why abstractions as we know them from the natural sciences do not occur in linguistics. It also explains why attempts at abstraction result in idealisations, with all the consequences we have outlined above.

If this is right, or at least in the right direction, as we believe it is, it has important consequences for the nature and the goals of theories in linguistics. More about that in Section 11.

10. Some more examples of idealisations in linguistics

The competence – performance distinction is by far not the only example of an idealisation by means of which modern linguistics has defined itself as a scientific discipline. Many of the consequences indicated in Section 3 appear to have characteristic features of idealisation, and not those of abstractions.

The idea of language as an infinite object, for example, is closely related to the competence – performance distinction. Modelled after concepts from the formal sciences (mathematical logic, mathematics, computer science)²⁵ this idea is based on the assumption that actual limitations on the use of, for example, embedding constructions (in terms of memory limitations, finite computational resources, and so on) are not intrinsically part and parcel of what language is. So what can be observed in reality, viz., that such limitations exist, is not considered to be an actual feature of the object ‘language’, but is taken to be ‘merely’ the result of intervening factors that as such are not intrinsically tied to the object. Language as we can observe it in actual use (in production and in interpretation) is a phenomenon in which unlimited recursion does not occur. Yet, it is being transformed into an ontologically different kind of object, for which there is no limit to recursive processes.²⁶

²⁴ And more generally, many of the humanities and social sciences.

²⁵ Cf., the afore-mentioned book by Tomalin (Tomalin, 2006).

²⁶ Cf., Fitz (2009) for extensive discussion, and a neural net model that is able to learn limitations of embedding constructions without an appeal to recursion. Cf., also Pullum & Scholz (2005).

Another example that was already mentioned is the characteristic, and almost exclusive, emphasis on written language. This also relates to the modelling of the object of modern linguistics on concepts from the formal sciences.²⁷ From a certain perspective the emphasis on written language seems quite justified: from a practical point of view written language is an object that is much easier to deal with than spoken language. Before the advent of sound registration equipment, writing was the only tool that could be used to collect speech and to share observations and analyses of it. In that light, traditional grammars can be considered as compact, codified reports on what could be observed in the field: speech. This is clearly a non-ideological, practical use of a methodological constraint. However, in modern linguistics such practical considerations are clearly not the only, or even the most important ones. The emphasis on written language also serves to treat ‘language’ as a well-defined, clearly delineated object. Speech is momentary, context-dependent, and seldom comes alone: prosody, gestures, facial expression, simultaneous interactions with elements of the non-linguistic context, it all occurs and happens at the moment of speaking, and that makes it difficult²⁸ to isolate as an object of study. Of course we can distinguish between sound and other components, but in particular when questions of meaning and interpretation are at stake that is in many cases not the relevant distinction: all components may contribute to the determination of what is being said. Hence, in so far as written language simply ignores these components, the transition from language as speech to language as writing is a clear example of an idealisation.

Along with this idealisation come yet others. One of them is the idea that competent language users can be considered as ‘disembodied’ individuals. Of course, embodiment is an essential property of human subjects, and moreover one that is in many respects connected with their being linguistic creatures. The body not only is an important intermediary with our physical environment,²⁹ it also plays a crucial role in determining the contents of large parts of our mental vocabulary, and it is a reservoir of all kinds of knowledge and abilities that are an integral part both of our linguistic competence, and of the way in which, and the ends to which, we use language. But neglecting embodiment

²⁷ Cf., e.g., Harris (2000) and the already mentioned Kraak (2008) for extensive discussion.

²⁸ And according to some even impossible; cf., Wittgenstein’s concept of a ‘language game’ that explicitly united both verbal and nonverbal elements (Wittgenstein, 1967, Section 7).

²⁹ Something that is revealed in language in many ways, for example in spatial indexicality.

has more effects than leaving out these essential features. The idealised competent language user whose linguistic competence is the central object of study of modern linguistics, is not just accidentally a disembodied subject, it is principally without a body. As was already noticed above, nothing in a theory about the resulting entity contains any lead as to how embodiment might be ‘added’ to it: the theory about the competent, disembodied language user is supposed to be a complete theory of human linguistic competence. From that perspective embodiment is not some real phenomenon from which we abstract, but an irrelevant property of human subjects.

A last example concerns semantics and pragmatics as branches of linguistic theory, and the central role played by the concept of ‘propositional content’. The dominant paradigm here relies on a principled distinction between propositional content as semantic meaning, and the use of expressions with such contents that results in pragmatic meaning. With the distinction comes a hierarchical relation: propositional content is independent from pragmatic meaning, whereas the latter needs the former as the base from which it is derived. This is the Gricean model and certainly within linguistics it is still the one used most.³⁰ In the philosophical literature the distinction as such has been subject of some debate.³¹ However, what is relevant to note here is that ‘radical contextualism’, the view that rejects the distinction, does not seem compatible with the goals of modern linguistics. And that indicates that the concept of propositional meaning as such is yet another example of a construction that is not so much an abstraction as an idealisation.

11. Consequences of idealisation

We hope that the foregoing discussion has made clear that the relationship between, on the one hand, the objects of study that modern linguistics has constructed via idealisations, and, on the other hand, language and linguistic competence as everyday, observable phenomena, is a complicated one, to say the

³⁰ Of course there are different views on what exactly the propositional content of an expression is, on how it is to be determined, and, consequently, where exactly the dividing line between semantics and pragmatics is to be drawn. But those discussions still operate within the assumption that the distinction, and the hierarchical relation between the two concepts of meaning, make sense.

³¹ Cf., various contributions in Preyer & Peter (2007).

least. This is something that Chomsky seems to recognise as well, as the following passage from Chomsky (1995, p. 20) shows:

At the conceptual-intentional interface [between sound-meaning pairs of *I*-language and actual language use] the problems are even more obscure, and may well fall beyond human naturalistic inquiry in crucial respects.

The construction of competence and the accompanying concept of an *I*-language (roughly, the ‘internal language’ which the idealised competent language user ‘cognises’) has distanced the object that according to Chomsky is the proper object of study so far from everyday language and its users that, as he himself acknowledges in this passage, it is not even clear which problems need to be solved in order for us to be able to relate them again. To put it differently, not only is there no bridging theory, it is not even clear what that theory is supposed to do. No doubt this aporetic situation is a direct consequence of the fact that it is not clear at all whether the idealised object puts any empirical constraints on such a theory, and if it does, what these might be.

To those who are primarily interested in language as an empirical phenomenon, Chomsky’s conclusion will no doubt sound quite defeatist. But Chomsky sees things differently. That, too, is clear from the passage just quoted: it is shown by his use of the qualification ‘naturalistic’.³² The use of this term reveals both a background ideology and an escape from this apparent impasse that Chomsky deems possible. What Chomsky aims at is not just some theory of language and linguistic competence, but one that is naturalistic through-and-through. Language and linguistic competence, as Chomsky sees it, are purely natural phenomena, of the same stature and nature as other human biological capacities and phenomena. For Chomsky the notorious claim ‘Language is an organ’ is not a metaphor (useful or not), but a factual statement. In the same way, and for the same reasons, that we study the human perceptual apparatus, the human motor system, and other biological capacities, with the means of the natural sciences, we cannot but study human linguistic competence, and hence human language, in the same manner.

So what motivates the idealisations Chomsky defends, is, as we have indicated before, an ideological position with regard to science and scientific

³² As for the qualification ‘human’ in ‘human naturalistic inquiry’: we don’t need to take that too seriously, we think.

method. It is scientific naturalism, plain and simple.³³ That such a choice for a strictly naturalistic methodology actually brings about a fundamental shift in ontology, is a consequence that Chomsky is apparently willing to accept, as the following passage from Hauser et al. (2002, p. 1570) shows:

The word ‘language’ has highly divergent meaning in different contexts and disciplines. In informal usage, a language is understood as a culturally specific communication system [. . .]. In the varieties of modern linguistics that concern us here, the term ‘language’ is used quite differently to refer to *an internal component of the mind|brain* [. . .]. We assume that *this* is the primary object of interest for the study of the evolution and function of the language faculty. [emphasis added]

But we do well to note that in this passage more is at stake than accepting the consequence that a naturalistic approach of language and linguistic competence studies a different object than another, more humanities-based approach. Apparently, the point is not to state that there are two (or more) alternative methodologies that we can choose from (and that we perhaps may provide arguments for a particular choice). Rather, what is claimed is that there is only one scientific approach possible in the first place, viz., the naturalistic one. Language and linguistic competence as they present themselves to us in real life, in observations about actual language use, simply are not phenomena that qualify for a scientific investigation.³⁴

A last observation concerning the position that is defended here by Chomsky and his associates concerns the scope of the resulting theory. That the linguistic competence of humans is rooted also in aspects of their biology is something no-one would doubt. That is a minimal rejection of an ontological dualism that seems quite generally accepted. The real question whether a theory that *reduces* the relevant core concept to biological entities and that accepts *only* a naturalistic methodology, will be able to come up with insightful explanation of properties of the original object of study. As the passage just quoted also illustrates, that seems to be a goal that Chomsky c.s. apparently are not willing to give up on. Their concern is ‘the study of the evolution and *function* of the language faculty’ [emphasis added]. Despite the pessimism that Chomsky dis-

³³ Cf., also Lappin et al. (2000) on this issue.

³⁴ It is also interesting to note that in this passage the authors speak of ‘the varieties of modern linguistics that concern us here’. Apparently, the present-day diversity of approaches (cf., Section 2) is something that the authors do acknowledge, if only by stating that alternative approaches do not ‘concern’ them. Cf., also Footnote 14.

played in the earlier cited passage from Chomsky (1995), the ambition to account for the function of language has not been abandoned, it seems. But in view of the ontological rift that the idealisations that are used have created, it certainly appears doubtful that this ambition can be realised.

12. Consequences of these consequences

What are the consequences for linguistics when its object of study is constructed via idealisation? Of course, it is not possible to answer this question fully and definitively. But what is clear is that the approach that modern linguistics has pursued over the last decades runs into a number of serious difficulties, difficulties that by the way also provide a partial explanation for the curiously diversified state in which linguistics finds itself today (cf., Section 2).

As was argued in the above in some detail, idealisation results in an ontological shift and creates an additional epistemological task, viz., the formulation of an adequate bridging theory. This leads to a number of problems. First of all, empirically motivated adequacy criteria for the bridging theory are very hard to come by: the idealised object itself does not deliver them, and observations with regard to the original phenomenon cannot function as such without further ado. This is a characteristic feature of idealisation, since, as we have seen, abstraction does not run into this problem. The second problem, which is an immediate consequence of the first one, is that there is a serious lack of empirical validation of the theory about the idealised object. Apart from the fact in the case of linguistics the original phenomenon is hard to fit into an experimental design, there is the problem that, without an independently verified bridging theory, no theory about the idealised object will lead to predictions that can be tested on the original phenomenon (via observation or by other means). And thirdly, as a result of that, the intuitive plausibility of the theory is seriously hampered.

Looking at the state of the art in applied linguistics, we see the consequences of this problematic situation clearly emerging. As the theoretical models of the generative tradition, based as they are on the notion of a grammar as a system of explicit rules, failed to deliver in applications such as machine translation, question-answer systems, and the like, people started to use other constructions of central concepts such as ‘language’, ‘meaning’, and so on. Often these new constructions were based on stochastic properties and patterns derived

from large corpora of actual text (and, later, speech). These constructions were based on other, often less far-reaching idealisations, i.e., they stayed closer to the original phenomenon and hence were more amenable to empirical testing. This development, however, is clearly motivated and steered by practical, pragmatic considerations, rather than by theoretical and explanatory ones. Theory, so it appears, lags behind practical application, which is also why we can observe a certain proliferation of theoretical models that are strongly influenced by very concrete, often also quite limited practical applications. To that extent, we might say that theory has become *ad hoc*.

In other contexts where linguistics touches on empirical research, another trend is visible. Language and linguistic competence are also important objects of study in the rapidly developing cognitive neurosciences. Inspired by a long tradition of psycholinguistic research, in particular research on language disorders, linguists have taken up the challenge provided by new, non-invasive techniques of studying the brain. The problems that occur here are partly related to the strongly naturalistic and reductionistic nature of a lot of neuro-physiological and brain research, partly they are due to the inherent limitations of the kind of experiments that the new techniques allow. One of the consequences is a reinforcement of the kind of idealisation that we have discussed in the foregoing, in particular the individualistic nature of the competent language user, and an accompanying diminishing possibility of linguistic theory to come up with leading hypotheses and testable predictions.

13. Conclusions

What conclusions can be drawn from these observations? Obviously, more research into the way in which linguistics, especially in its present-day diversity, copes with its central concepts, is needed. But one question will be central: Is naturalism in linguistics a methodology that is forced upon us by the nature of the phenomena it studies? Or is it a choice? The observations and considerations put forward in this paper strongly suggest that the latter answer is the correct one: the naturalism that is so characteristic for modern linguistics, in particular, but not exclusively, for the generative tradition, is based on a scientific ideology. Note that as such, that does not imply that the resulting methodology is necessarily the wrong one. (It could be the right choice made for the wrong reasons.) But it does show along which lines further research in

this area should be conducted: it is the consequences of this choice that need to be thoroughly scrutinised.

Should it turn out, as we strongly suspect it will, that the ideologically motivated choice for naturalism severely hampers the explanatory power of the resulting linguistic theory, then that by itself provides a clear pointer to the direction in which one may look for alternatives. For that a naturalistic approach that is not ideologically motivated may lead to interesting and, to some extent, testable results is shown by various alternative theoretical frameworks that, partly as a response to the deficiencies of work done in the generative tradition, have been developed over the last decade or so. Examples are cognitive linguistics³⁵, stochastic linguistics,³⁶ and approaches in which neuronal models of language acquisition and language use are studied.³⁷

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

References

- Chomsky, Noam. 1965. *Aspects of the theory of syntax*. Cambridge MA: MIT Press.
- Chomsky, Noam. 1980. *Rules and representations*. Oxford: Blackwell.
- Chomsky, Noam. 1986. *Knowledge of language*. New York: Praeger.
- Chomsky, Noam. 1995. Language and nature. *Mind* 104. 1–61.
- Daelemans, Walter & Antal van den Bosch. 2005. *Memory-based language processing. Studies in natural language processing*. Oxford: O.U.P.
- Fitz, Hartmut. 2009. *Neural syntax*. Amsterdam: ILLC/Department of Philosophy dissertation.
- Groenendijk, Jeroen & Martin Stokhof. 2005. Why compositionality? In Greg Carlson & Jeff Pelletier (eds.), *Reference and quantification: The Partee effect*, 83–106. Stanford CA: CSLI-Publications.
- Harris, Randy Allen. 1993. *The linguistics wars*. Oxford: O.U.P.
- Harris, Roy. 2000. *Rethinking writing*. London: Athlone.
- Hauser, Marc D., Noam Chomsky & Tecumseh W. Fitch. 2002. The faculty of language. *Science* 298. 1569–1579.
- Jackendoff, Ray. 1987. *Consciousness and the computational mind*. Cambridge MA: MIT Press.

³⁵ Cf., the recent Tomasello (2003).

³⁶ For an interesting recent example, cf., Daelemans & Van den Bosch (2005).

³⁷ Cf., e.g., MacWhinney & Chang (1995) and the afore-mentioned Fitz (2009).

- Jones, Martin R. 2005. Idealisation and abstraction: A framework. In Martin R. Jones & Nancy Cartwright (eds.), *Idealization XII: Correcting the model. Idealization and Abstraction in the Sciences. Poznań studies in the philosophy of the sciences and the humanities*, 173–217. Amsterdam: Rodopi.
- Kamp, Hans & Martin Stokhof. 2008. Information in natural language. In Johan van Benthem & Pieter Adriaans (eds.), *Handbook of philosophy of information*, 49–112. Amsterdam: Elsevier.
- Kraak, Albert. 2008. *Homo loquens en homo scribens*. Amsterdam: Amsterdam University Press.
- Lappin, Shalom, Robert D. Levine & David E. Johnson. 2000. The structure of unscientific revolutions. *Natural Language and Linguistic Theory* 18. 665–671.
- MacWhinney, Brian & Franklin Chang. 1995. Connectionism and language learning. In Charles A. Nelson (ed.), *Basic and applied perspectives on learning, cognition, and development*, 33–58. Abingdon: Erlbaum.
- Newmeyer, Frederick J. 1986. *The politics of linguistics*. Chicago: The University of Chicago Press.
- Postal, Paul M. 2004. *Skeptical linguistic essays*. Oxford: O.U.P.
- Preyer, Gerhard & Georg Peter (eds.). 2007. *Context-sensitivity and semantic minimalism. New essays on semantics and pragmatics*. Oxford: O.U.P.
- Pullum, Geoffrey K. & Barbara C. Scholz. 2005. Contrasting applications of logic in natural language syntactic description. In Peter Hájek, Luis Valdés-Villanueva & Dag Westerståhl (eds.), *Logic, methodology and philosophy of science: Proceedings of the twelfth international congress*, 481–503. London: King's College Publications.
- Stein, Harry. 1997. *The fibre and the fabric*. Amsterdam: ILLC/Faculteit der Wijsbegeerte dissertation.
- Stokhof, Martin. 2002. Meaning, interpretation and semantics. In Dave Barker-Plummer et al. (eds.), *Words, proofs and diagrams*, 217–240. Stanford CA: CSLI-Publications.
- Stokhof, Martin. 2007. Hand or hammer? On formal and natural languages in semantics. *Journal of Indian Philosophy* 35. 597–626.
- Stokhof, Martin. 2008. The architecture of meaning: Wittgenstein's *Tractatus* and formal semantics. In David Levy & Eduardo Zamuner (eds.), *Wittgenstein's enduring arguments*, 211–244. London: Routledge.
- Tomalin, Marcus. 2006. *Linguistics and the formal sciences*. Cambridge: C.U.P.
- Tomasello, Michael. 2003. *Constructing a language: A usage-based theory of language acquisition*. Cambridge MA: Harvard University Press.
- Wittgenstein, Ludwig. 1967. *Philosophische Untersuchungen*. Frankfurt/M.: Suhrkamp.