What Cost Naturalism?

Martin Stokhof & Michiel van Lambalgen[*]

The paper traces some of the assumptions that have informed conservative naturalism in linguistic theory, critically examines their justification, and proposes a more liberal alternative.

1   Introduction

In this paper we take up an issue that was touched upon in our earlier paper on abstraction and idealisation in linguistics[1] only in passing, viz., what we there called 'the ideological nature' of certain views about the nature of linguistics as a scientific enterprise. The choice of the term 'ideological' has confused some readers and may have been less fortunate. But apart from the choice of appropriate terminology, there is, we feel, an important issue here, one that needs further investigation. This note is a first step.

What is the issue? To put it in general terms, many approaches in modern linguistics are characterised by an, implicit or explicit, commitment to a concept of language that views it as the kind of natural phenomenon that can be studied by scientific means, with the natural sciences (physics, biology) acting as role models for what proper scientific theorising looks like. Language here is an ontologically homogeneous phenomenon that, in principle, can be captured and explained completely in, broadly speaking, physicalistic terms.

The original remark about the ideological nature of this view is motivated by our concern that it is insufficiently based on a prior and independent conceptualisation of what language is and what an explanatory theory of it would need to account for,

---

[1]   Stokhof & van Lambalgen (2011a).

Kata Balogh and Wiebke Petersen (eds.). 2017. *Bridging Formal and Conceptual Semantics. Selected papers of BRIDGE-14.* Düsseldorf: dup.

91

with only subsequently an argumentation that such goals can actually be achieved best by a physicalistic theory. Rather, it seems, the reverse has taken place. With the choice of a particular type of theory already in place, the concept of language has been adjusted and changed so as to fit the pre-conceived idea of what a proper theory of language should look like. But such a move can only be inspired by the idea that only the type of theoretical explanation that we know from the natural sciences can count as a proper account of whatever phenomenon we are dealing with. And that, we venture to claim, is not motivated by argument but by ideology.

   In order to make good on this claim, we need to do a number things. First of all, we need to provide evidence that this kind of reasoning is indeed used in discussion about what the nature of linguistics is. Second, we need to show that this leads to the ontological homogenisation that we claim it does. That should settle the 'ideology' claim. But if, thirdly, we also are able to show that the resulting picture is deficient, both descriptively and explanatorily as well as philosophically, our investigation will also be able to provide some support for alternative conceptions of language and for other ways of doing linguistics.

2   A quick exploratory dig

One prevalent view on naturalism as applied to linguistics, to which we will turn in the next section, is, we venture, a reflection of what one might call 'Chomsky's shadow', a result of the deep and still present influence of some key assumptions that were infused in modern linguistics with the advent of generative grammar. There are many such elements that continue to shape theoretical thought in the discipline, here we focus on two that we think are directly relevant for the topic of this paper, viz., individualism and the adoption of a particular type of explanation. We realise of course that what follows is only a very rough sketch, one that traces one particular way of thinking, and that does not do justice to some of the alternative approaches that have been developed.[2] Nevertheless, we do think that the underlying assumptions that we discuss are not operative only in one particular paradigm, but that they are around in other frameworks as well.[3]

---

[2] Such as various approaches in functional linguistics and cognitive linguistics. Cf., e.g., Givon (2013) for some discussion of the assumptions underlying functional linguistics that is congenial with the argument developed in this paper.

[3] Cf., below, footnote 13 and 22 for two examples from different frameworks.

The individualism that is part and parcel of many approaches in linguistics, past as well as present, can be stated succinctly as follows: 'language is an individual asset, and linguistic ability is an individual property'. The assumption is that in principle, though not de facto, an individual could be a competent language user all by his- or herself. Or to put it slightly differently, in giving a description of what linguistic ability is, and in giving a description of what that ability is about, viz., language, there is no need to refer to anything over and above the individual itself. Language use may be a social art, but language itself and the ability to use it are not. Of course, in real life people do not become language users in complete isolation, nor do they function as such without being part of a social environment. But, or so the idea goes, 'in principle' these social aspects can be bracketed, at least if they themselves are not a topic of study, of course. With regard to individualism, then, the basic divide is not between theories that leave out the social (communicative) dimension as a proper concern for linguistics and theories that include it, but between those that assume that the social dimension can be viewed in terms of the interplay of individual abilities and those that do not view the individual ability as an independent and foundational element.

The second factor that we want to draw attention to is the reliance on a particular form of explanation. There is a tendency in many linguistic theories to assume that explanations of linguistics facts ultimately need to be stated in terms of structural properties, of language, grammar, and of individual language users.[4] This inspires an exclusive focus on structural properties of the human cognitive make-up, ultimately, those of the human brain. This notion of a structural explanation is familiar from the sciences. However, in the case of linguistics actual access to such underlying structural properties never was a substantial part of the enterprise: it remained, rather, an assumption that such access would be possible, i.e., it was a way of formulating explanations, rather than a way of explaining things. This is what Ernan McMullin called 'hypothetico-structural explanation (McMullin 1978, p. 139):

> When the properties or behaviour of a complex entity are explained by alluding
> to the structure of that entity, the resultant explanation may be called a structural
> one. […] Such explanations play only a small role in scientific enquiry. Much

---

[4] This is not to say that the resulting theoretical frameworks will be the same: what counts as the relevant structural properties, what form the system takes, what role semantics and pragmatics have to play in an overall account, are some of the parameters along which different approaches distinguish themselves.

commoner are those where the structure is postulated to account for the observed properties or behaviour of the entity under investigation. [...] [These] could be called 'hypothetico-structural'.

When individualism and the concept of hypothetico-structural explanation are combined the distinction between competence and performance that has been enormously influential, and not just in the generative grammar tradition,[5] seems almost self-evident. And with that the first, decisive step on an ever more abstract construction of the central object of linguistics is taken that we know from the generative tradition: from actual languages, as used in the 'here-and-now', to the concept of 'possible human language', and then to that of 'universal grammar' and ultimately the 'faculty of language in the narrow sense'.

Of course, not all linguistic theories have been travelling that far on the road of abstraction, but many of them have been, and still are, affected by this powerful combination of assumptions.

3   Abstraction and idealisation, once more

Now, one could regard the movement just described also in a positive way, i.e., as a manifestation of the increasing maturity of linguistics as a scientific enterprise. And it has been presented as such in the literature. After all, no scientific discipline engages directly with the continuous stream of experiences of phenomena that we have, they all construct their object of investigation from the endless particularities of phenomena by focussing on what is deemed important and leaving out the rest. Thus any discipline needs to create from the experiental flux a more or less stable set of entities that it aims to investigate.

That much is certainly true, and it is something can be observed in any scientific discipline (as well as in many other branches of human cognitive activity, by the way.) But as we have argued elsewhere,[6] there are crucially different ways of doing this, that have substantially different implications for the relationship between what a discipline is concerned with and the phenomena themselves.

Very briefly, in the paper just referred to, we made a distinction between 'abstraction' and 'idealisation', and characterised the difference in broad terms as follows.

---

5   It was also a formative element in the development of formal semantics, for example.
6   Cf., Stokhof & van Lambalgen (2011a, 2011b).

Features of a phenomenon that are abstracted are real features that at some point in time are considered to be too complex or too intractable, or, in some cases, not sufficiently relevant, to be taken into account in conducting a scientific inquiry into the nature of the phenomenon. A decision to abstract away from a feature is thus context-dependent and reflects various types of constraints that may obtain at a particular moment in time, relating to the availability and accuracy of instrumentation, availability and access to data, and so on. What needs to be pointed out is that abstraction is an intentional move: the features that are abstracted from are acknowledged as real, and they do occur, albeit in a special way, in subsequent theories, and, being actual features of the phenomenon, they will manifest themselves in experiment and observation. That holds, of course, also for features of a phenomenon that have not been acknowledged as such, i.e., features that have not been observed. These play a role in experimental and observational results as well, but they are not intentionally left out of the theory. They are simply not taken into account because they have not been observed.

Features that are abstracted from are typically quantitative in nature, and often they concern the numerical value of something that is known to exist, but difficult, or in some cases not particularly relevant, to measure precisely. Examples for the natural sciences would be movement on a frictionless plane; the concept of a perfectly rigid rod, or of perfect vacuum; free space constants, such as the gravitational and magnetic constants; or the concept of a perfectly pure chemical substance. In each of these cases there is a quantitative parameter (friction, the number of molecules in a certain volume, etc) that is set to a particular value (zero, infinity, or a specific number, as the case may be), not because that value is known, which it isn't, but because it is too difficult or too unimportant to actually measure.

What should be kept in mind is that although abstracted from, these parameters not only do turn up in experiment and observation they also are essential ingredients of the theories that are based on them. If we described the movement of an object on a frictionless plane, we are not assuming that friction does not exist; we're only setting its value to zero. That is of crucial importance because it implies that there will always be discrepancies between the predictions made by a theory that is based on the abstraction and what we can observe in the laboratory or in the real world. And it are such discrepancies as we would be able to observe and eventually measure, e.g., when better or new instruments become available, that will allow us to provide richer and better theories that rely on less abstractions.

Thus in the case of abstraction there is always a 'back-and-forth' between the theory based on an abstraction and the real phenomenon as it manifests itself in experiment and observation. And it is this back-and-forth that allows, and, in some cases, forces us to 'undo' the abstraction and come up with a better theory.

So the motivation for abstraction is primarily methodological (in a fairly literal sense, as 'having to do with what methods we have at our disposal') and practical, and hence in principle always temporary and revisable.

What about the second way of constructing an object of investigation, viz., 'idealisation'? The difference with abstraction is that here features that are 'idealised away' literally disappear from view. These are features of the phenomenon that are considered to be irrelevant from the perspective of what the discipline is interested it. Hence an idealised feature no longer plays a role in the empirical investigation: it is declared unnecessary to be considered further. The difference with abstraction may be subtle, but is it fundamental: an idealised feature is not 'merely' too complex or intractable, or not relevant in a specific use case, rather, it is viewed as something that need not be considered in, and may even stand in the way of, an attempt to understand the phenomenon since it is irrelevant for acquiring such an understanding.

Thus idealisation creates an ontological gap between phenomenon and constructed object: an idealised feature is, of course, a real one at the level of the phenomenon, but it is no longer present in the constructed object and has no role to play in the subsequent theories that employ that object. In a quite literal sense idealisation is an ontological move, rather than an epistemological one like abstraction: it changes the subject.

Features that are idealised are typically qualitative features of the phenomenon. Examples from linguistics that suggest themselves are: language as an infinite set of sentences; the competence – performance distinction; or the concept of literal meaning. In each case there is a particular feature of actual language use (e.g., its dependence on non-linguistics aspects of context, or the limitations of cognitive processing power, and so on) that is being discarded. As a result a new concept is formed that lacks this feature, and thus constitutes a different kind of entity than the original phenomenon.

Take the first example: if we look at languages in the 'here-and-now', i.e., as the phenomenon that we actually encounter when people use language, it is obvious that it is a finite object. The total number of utterances of any natural language,

past, present and future, is finite, as is the existence of the human species. What is important is that this finite number of utterances appears to be unpredictable in this sense that we constantly encounter utterances of new expressions and that we have no reason to assume that this is bounded purely quantitatively (i.e., there there is some fixed n that serves as the upper bound of the number of utterances we may encounter). Thus 'creativity' of language use/users is a real phenomenon, one that needs to be accounted for. What is important to note is that this creativity in fact is bounded: it does not mean that anything goes and that there are no restrictions, e.g., on the length or the complexity of the structure of utterances (especially embedding) that are operative

And this is were things go wrong. The preferred way of accounting for creativity is to drop the actual feature of language's finiteness and switch to a different concept of language that identifies a language with an infinite number of expressions.[7] But this is an ontological switch that is underdetermined by the observed creativity, and that turns a fact to be explained into a defining feature, and thereby changes the phenomenon that we observe into a different type of entity altogether.[8]

This also means that with idealisations there is no straightforward 'back-and-forth' between the theory containing the idealisation and the actual phenomenon as it appears in experiment and observation. Since the idealised feature is missing from the theory, there is not direct relation between what we learn from experiment and observation and the predictions made by the theory. There is an ontological gap, and that gap needs to be bridged by an additional 'bridging theory' that relates what the theory says to what is actually out there. That means that idealisation comes with an additional epistemological task, viz., to come up with an adequate bridging theory. How to do that is not at all obvious, however, especially because the adequacy criteria for such a theory seem hard to come by.

So unlike abstraction, which is clearly motivated by practical and methodological concerns, the motivation for idealisation is different: it does not derive from empirical considerations, but rather is motivated by conceptual-philosophical ones. Where these are based on pre-conceived ideas about what constitutes a proper scientific investigation, –and as we shall argue, such ideas are indeed

---

[7] A move that is inspired by the concept of a language that comes from the deductive sciences.
[8] For another example, viz., the celebrated 'competence – performance' distinction, cf., Stokhof & van Lambalgen (2011a, section 8).

behind idealisations that we find in linguistics,– they can be properly be called 'ideological'.

4 Naturalism in linguistics: a specific case

Before turning to a more general outline of naturalism in linguistics, we first briefly discuss a specific way of arguing for this general view. It draws on the work of Chomsky, arguably one of the most influential thinkers in this area, and not just in the generative tradition.[9] Chomsky has argued extensively that linguistics is a science like other sciences, and hence that it should follow the same leads in the construction of its objects as well as in the methodologies it employs in studying them. That means not only that linguistics should be held to the same standards as any other empirical discipline, but also that it should be judged by similar criteria. This, Chomsky claims, is not always the case, especially in philosophy, where what is accepted as normal and appropriate in the other sciences is judged by other standards if it occurs in linguistics (Chomsky 1995, p. 7):

> [...] it is a rare philosopher who would scoff at its [i.e., physics', MS-MVL] weird and counterintuitive 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.

The message here is clear. According to Chomsky philosophers accept what is done in physics and related areas of research at face value, and they study it in order to understand it the way it is. But when it comes to linguistics and similar disciplines the subject matter is not studied 'as is', but it is judged, and apparently by criteria that are germane to the disciplines in question.

---

[9] The 'present relevance' of Chomsky's thought is a matter of debate. It is certainly true that the landscape of theoretical linguistics is much more heterogeneous now than it was in the sixties, seventies and eighties of the previous century. And many approaches are self-proclaimed 'non-' or even 'anti-'Chomskyean. However, there is still a strong tradition in theoretical linguistics that subscribes to some of the fundamental principles discussed here. Cf., below for some examples. And we would claim that even in alternative theoretical approaches one can find assumptions that derive from the generative tradition. For more discussion and concrete examples, cf., Stokhof & van Lambalgen (2011b, pp. 80–85).

But clear as it is, what Chomsky claims here is not beyond dispute. There is a lot of critical discussion of various aspects of physics, or at least of the conceptual understanding of it, both in philosophy as well as in physics itself. The ongoing debate on the proper interpretation of quantum mechanics immediately comes to mind (and we will see a bit of that later on), but also more generally there is critical engagement with the natural sciences and their self-understanding.[10] So it is definitely not true that philosophers do not criticise basic concepts and fundamental principles in the sciences.

Likewise, the charge brought against philosophy that it judges linguistics by 'outside' criteria does not seem completely justified either. Of course, as in any area of investigation, here too there is no doubt that examples can be found of analyses and criticisms that are 'off the mark'. But that is not particularly interesting. Rather, what Chomsky seems to take issue with is criticisms that start from other assumptions regarding what is the proper methodology in linguistics than he deems relevant, and that hence accept methodological pluralism. This is mocked by Chomsky as follows (Chomsky 2000, p. 76):

> We must abandon scientific rationality when we study humans 'above the neck' (metaphorically speaking), becoming mystics in this unique domain, imposing arbitrary stipulations and a priori demands of a sort that would never be contemplated in the sciences.

It is here that a central element of Chomsky's particular form of naturalism comes to the fore. As the passage just quoted strongly suggests, there is, according to him, no special methodology that needs to be employed when studying the human mind, including language. Such an assumption can only lead to 'mysticism', not to explanatory theories. Rather, there is every reason to extend the methodology employed in the study of humans as physical (biological) entities to the study of mind and language.

Thus it appears that Chomsky's main gripe with his philosophical opponents is that they do not share his assumption that, linguistics being an empirical discipline, it should employ the same methodology as any other empirical science, i.e., that they do not, as he does, subscribe to methodological monism.[11]

---

[10] A nice example is Hartry Field's discussion of the role of mathematics in physics (Field 1980); an example in biology is the debate concerning the unit of selection in evolution.

[11] That Chomsky's ideas are still with us is testified by the way in which the nature and goals of linguistic theory are described in more recent work. Here is an example from Culicover and Jackendoff (Culicover & Jackendoff 2005):

5   Naturalism in linguistics: three characteristics

The methodological monism that we identified in the previous section as the rock-bottom of Chomsky's distinct view on the nature of linguistics, is an important characteristic of a naturalistic stance in linguistics (as it is in many other disciples). In this section we briefly review three distinct such characteristics. Of course, in real life naturalism comes in many different forms and guises, and with different motivations and justifications. However, we do claim that these three features are characteristic in the sense that they can be found, perhaps in mixed composition, i.e., with different emphasis and motivation, in most naturalistic views.

The three characteristics are methodological monism, already introduced; ontological monism; and explanatory monism. The first is the assumption that all of science employs the same methodology. This should be read in a broad sense, since, obviously, not all concrete methods are equally relevant, or even applicable in every discipline. The second characteristic embodies the conviction that there are no major distinct ontological categories, i.e., that everything is basically made of the same 'stuff'. Again, this is a claim that should be interpreted with some caution, it is not necessarily a denial of any form of categorial distinction between entities, but embodies the claim that all of them are part of the same ontology. The third characteristic, finally, comes with the claim that there is one model of explanation that fits all phenomena, i.e., that all scientific disciplines employ the same requirements for what constitutes an adequate explanation of the phenomena that constitute their subject matter.

The relations between these three characteristics are complex. It is important to note that they are not merely three sides of one and the same (thick) coin, i.e., they do not imply one another. For example, one could subscribe to ontological monism,

---

We begin a more thorough examination of the situation [in syntax, MS-MVL] by reviewing the first principles of generative grammar, articulated in detail by Noam Chomsky in *Aspects of the Theory of Syntax* (1965) and many subsequent works. With only minor modulation and reinterpretation, these principles have stood the test of time and have received further confirmation through the flood of research in cognitive science in the past forty years. [...] Generative grammar is grounded in the stance that the object of study is the instantiation of language in the context of the human mind/brain, rather than an abstract phenomenon that exists 'in the community' [...] The fundamental linguistic phenomenon is a speaker producing an utterance that is understood by a hearer, and the fundamental question is what is present in the speaker's and hearer's mind/brain that enables this interchanges to take place.

From there Culicover and Jackendoff proceed to identify 'productivity' and 'competence' as core concepts.

yet acknowledge distinctly different methodologies as required by different sets of ontologically non-distinct phenomena. What does follow is that in such a case the motivation for the methodological pluralism can not be ontological. But other justifications can be thought of.

Of course, despite the conceptual independence of these three forms of monism, there are in fact strong affinities between them, and people tend to embrace them as a kind of package deal. But their unification does require some additional assumptions. One such assumption that enables further identification of these three forms of monism is that nature is material. Ontological monism then entails that all entities are material, which in its turns lends much plausibility to methodological monism: if what we study is basically of the same kind, then obviously the same methodology applies across the board (of disciplines, not entities). A further assumption then comes into view, viz., that there is a basic level of material constitution at which all relevant phenomena can be studied. This reductionism then further motivates explanatory monism: if what we study with the same methodology takes place at the same level of material constitution, how could the explanations we are after be different?

Thus arises the idea of a unified science, based on material reductionism. With physics arguably being the most successful science we have, it would stand to reason that this unified science employ the methodology of physics, and if we accept that, we find ourselves engaged with some form of physicalism.

In order to make this a plausible account for linguistics, we actually also need as some form of (methodological) individualism as an additional assumption. This individualism was identified in section 2 as one of the core contributions that Chomsky has made to linguistic theory. It seems to work because the physical realisations of all entities appear to be individual in nature. That holds in general, and for linguistics it seems to imply that it is individual language users that are the core entities that linguistic theory should be concerned with. Thus the individualism appears to guarantee ontological homogeneity, and thereby it also seems to guarantee explanatory unification.

As we have argued elsewhere,[12] some strong idealisations, in particular the competence – performance distinction, are needed to make this work in the case of language. However, as pointed out above, the argumentation for these assumptions is conceptual, and not (or at least not exclusively) empirical. And that suggests

---

[12] Cf., Stokhof & van Lambalgen (2011a, section 8).

that at least in linguistics this form of naturalism is ideologically motivated. It is not based on empirical arguments for methodological monism, nor does it give a sound empirical argumentation for the necessity of methodological individualism. Rather, these two crucial assumptions are embraced on the basis of what ultimately is a scientistic ideology.[13]

6   Naturalism: an alternative view

At this point it might be good to pause for a moment and ask ourselves what this scientific form of naturalism has bought us in linguistics, and whether there might be an alternative way of being a naturalist and a linguist.

As for the first question, we would venture that the answer is: 'Actually, not that much.' It is true that linguistics enjoys considerable academic prestige, and in part (but admittedly, only in part) this is due because it conceives itself, and is conceived by others, as being the most scientific (rigorous, formal, …) of the humanities. The adoption of a naturalistic stance no doubt contributes to that. But in actual practice it has led, or so we would argue, initially to an unjustified restriction of the domain of inquiry, and later, when those restrictions became too stringent, to a confusing variety of approaches and methodologies. The continuing adherence to individualism and the model of hypothetico-structural explanation have played a major role here.

---

[13] To further counterbalance the impression that we are only after Chomsky, or generative grammar, we illustrate very briefly how similar ideas have crept into our own work in semantics. In our use of the event calculus, which uses closed world reasoning in the description and explanation of a range of phenomena, from tense and aspect to coercion and logical reasoning, some assumptions about individualism and materialism can be discerned. Cf., the following passage from Stenning & van Lambalgen (2008, p. 161–162):

> Systems of closed-world reasoning are logics of planning. […] Maintaining a model of the current state of the immediate environment relevant to action is a primitive biological function. […] The planning logics are just as much what one needs for planning low-level motor actions such as reaching and grasping, as they are for planning chess moves.

> Approaching from the direction of the syntactic and semantic analysis of temporal expressions of natural languages also directs attention to planning as underlying our faculties for language. More generally, a main human brain innovation is the increase in neocortex, and specifically in frontal areas of neocortex.These frontal areas are involved in planning and 'executive functions', among other things.

Clearly the assumptions that are operative here are that language originates from non-linguistic structural properties of the human brain; and that, hence, core aspects of linguistic meaning can be explained in terms of such individual structural properties.

But what about the second question, about the possibility of an alternative? If we bracket our commitment to individualism combined with hypothetic-structural explanation, the most intuitive way of explaining the human ability to use language would seem to be one that proceeds in terms of dispositional properties of language users. Instead of postulating 'underlying' structural properties that explain what we can observe about language and its use, — a postulate that would seem to bring along a mechanistic picture of language and language users —, we would try to explain what we can observe in terms of dispositions to verbal and relevant non-verbal behaviour that are triggered in a complex setting of natural, individual, and social conditions.

Such a dispositional account has a number of attractive features. First, it would provide an intuitive account of the voluntary nature of language use. This intentional aspect is hard to give a place in the structure-based, mechanistic view, but comes natural in a dispositional account. After all, dispositions need additional conditions to manifest themselves, and an intentional volitional act would appear to be a prime candidate for such a condition. (Which is not to say that it would be a necessary condition.)

Secondly, it would provide space to accommodate a number of factors other than the individual itself that play a role in a comprehensive account of language and language use. One example here would the communicative purposes of language users, that are relevant input for determining the manifestations of their dispositions to use language. That is still at the level of language users and their interactions. But the dispositional view could also accommodate other factors, such as external determinants of (lexical) meaning, or institutions and other social entities. Since the manifestation of a disposition is something that 'takes place', i.e., is a spatio-temporally located event, the role of both physical and social reality as (co)determining meaning can be incorporated in a fairly straightforward manner.

And thirdly, unlike the competence – performance distinction, which isolates competence from performance limitations, such as constraints on working memory, or limited attention span, the dispositional view prepares a natural setting for such limitation to operate in. After all, as noted above, the manifestations of a disposition take place in concrete situations, and features of these situations may not only contribute to the content of these manifestations, they may also set limits to them, in various ways.

So, all in all, it seems much more intuitive to look at what is individual about language and language use in terms of dispositions than in terms of structure, and to account for them as embedded in a rich and varied social ontology.

However, enter the commitment to hypothetico-structural explanation, again. It is this commitment, that postulates that, 'in the end', explanations will refer only to structural properties, that forces us to consider an explanation in dispositional terms as 'provisional', at best. If not outright non-explanatory,[14] then minimally these are explanations that contain terms, viz., the ones referring to dispositions, that will need to be analysed further until they are eliminated and only reference to structure and structural priorities remains.[15] And this is because, so the argument goes, structure is real and dispositions are not.

As was already mentioned, a key element in the justification for this assumption, and hence for the commitment to hypothetico-structural explanations, is the conviction that physics works that way. And given that physics is the most successful scientific discipline there is, the idea is that linguistics should follow its lead. This raises two questions. First of all, does physics really work that way? And secondly, and this is an issue that is quite independent from the answer to the previous question, what is the justification for linguistics to follow suit?[16]

The first question is a complicated one. The idea that structural explanations are the *nec plus ultra* of all scientific inquiry is based on what is often called the 'Groundedness Thesis', i.e., the view that all dispositions are ultimately grounded in structural properties of the entities involved. As said physics is often assumed to conform to this assumption, but there are dissenting voices. More specifically, it has been argued that at the level of quantum phenomena we find properties that are

---

[14] Arguments to that effect are, by the way, more often than not quite rhetorical in nature, and consist mainly of obligatory references to the circular nature of, e.g., an explanation of the effect of opium in terms of its 'dormative' power.

[15] Getting rid of the non-individual entities would be another task.

[16] It is good to note at this point that although naturalism is very much the default position among linguists in the generative tradition, there are exceptions. An example is provided by Hinzen and Uriagereka, who defend a non-naturalistic interpretation along the following lines (Hinzen & Uriagereka 2006, p 71–72):

> In our view the metaphysics of linguistics points to a radically different ontology of the mind that invites a rather novel reflection on the constraints that delimit the human conceptual edifice, and which to this day have no plausible biological or even physical explanation. As a consequence of that, the human faculty poses much the same explanatory problems for contemporary physicalism as the mathematical faculty does.

This is an unorthodox position that deserves further scrutiny. Unfortunately to do so is beyond the scope of this paper and has to be deferred to another occasion.

inherently dispositional, i.e., not reducible to structural properties that can serve as their causal basis. One example is spin of an electron. Cf., Bigaj (2012, p. 212):

> Orthodox quantum mechanics does not seem to identify any property which could play the role of the causal basis for spin.

But also position arguably needs to be considered as a disposition, rather than a classical property, in view of the probabilistic nature of quantum phenomena. In fact, this holds across the board, given the very nature of quantum phenomena (Bigaj 2012, p. 214)]:

> The dispositional interpretation of quantum properties is a direct consequence of the probabilistic character of quantum states.

Thus the defining characteristic of quantum theory, viz., it is inherently probabilistic character, enforces a dispositional view of its basic objects and their properties. And according to some this even applies to the very existence of quantum objects. Cf., Thompson (1988, p. 77):[17]

> In quantum field theory (a more complete form of quantum physics), even the existence of objects is a dispositional property that may or may not be manifested, as, for example, pairs of particles and anti-particles may or may not be formed.

So, taking the lead of science does not force one to do away with dispositions: even the most fundamental theory of physics does not obey the model of hypothetic-structural explanation, because, as one might put it, 'physical reality is dispositional'.

Now this view on physics and its explanatory model is not uncontroversial, and one might very well maintain that ultimately a fully explanatory theory of the physical world must satisfy the Groundedness Thesis. Which brings us to the second question raised above: Even if physics could rely solely on structural explanations, what reason is there to think that linguistics can too? In order to answer this question, we need to take a step back and look at what motivates naturalism in the first place.

7 The basic claim of naturalism

The core of the naturalistic stance, i.e., the assumption that different forms of naturalism all subscribe to, is that there are no supernatural entities that we

---

[17] Cf., also Suarez (2007) for some further discussion.

need to appeal to in the construction of adequate, explanatory theories of natural phenomena. Or, to put it in the form of a slogan: 'Nature can be explained on its own terms.'

The reasons for making this claim fall in two broad categories. One is ontological/metaphysical: we need not appeal to supernatural entities because such entities do not exist, i.e., they are not part of the ontological furniture of the world. The other is epistemological/methodological: an appeal to supernatural entities is not needed because such entities do not appear in our explanations, i.e., they do not belong to the ontological furniture of our theory of the world. These two types of motivation are not completely independent, of course.

But what counts as a supernatural entity? One way to define the concept would be to give a list of what it applies to, i.e., a list of supernatural entities that have been assumed, conceptualised, or otherwise 'identified' over the ages. It would contain not only well-known concepts from religion, magic and kindred realms, but also some that did occur in scientific explanations at some point in time: phlogiston, vital force, ... Such an enumeration suggests also a more general characterisation, which typically proceeds in terms of what a supernatural entity is *not*:

> x counts as a supernatural entity if and only if x lacks relevant features of physical entities, such as having material constitution, being situated in space-time, being subject to laws, …

To be sure, this is still a fairly general and 'open' characterisation, but it will do to illustrate that the driving force behind naturalism, at least initially, is to avoid ontological and/or epistemological excess. What does not share basic characteristics with what nature has on offer does not belong to nature but is (or rather: would be) of an entirely different order.

The reference to physical entities in this characterisation might suggest that naturalism and physicalism, which holds that all sciences can be reduced, in some sense, to physics, are actually two sides of the same coin, but that would be too hasty a conclusion. For note that being a naturalist comes with the commitment to eschew reference to supernatural entities, not with the much stricter injunction to refer only to physical entities. What makes its appearance in one's naturalistic explanations must have a physical realisation, but that falls short, or so we will argue, of being a physical entity.

8   Two flavours of naturalism

In order to bring out the difference we distinguish between two types of naturalism: conservative, and liberal.[18]

Conservative naturalism represents the more strict view that not only all natural entities need to have physical realisations, but that all entities *are* indeed physical entities. It is in that respect that liberal naturalism takes a different stance: it agrees with conservative naturalism that all entities need physical realisations, but nevertheless maintains that not all such entities are therefore physical entities, or can be reduced to such entities. In other words, liberal naturalism leaves open the possibility that there are natural entities that are not 100% physical in this sense that, although they do depend on physical entities that realise them, they can nevertheless not be reduced to their physical realisations.

It will be clear from the above that there are natural affinities between conservative naturalism and ontological monism, in particular physicalism, and, in its wake, with methodological monism and explanatory monism. The conservative naturalist position is indeed, as De Caro and MacArthur quite rightly note, a classical one. It is strongly committed to the 'unified science' ideal of logical positivism in a physicalistic form: there are no other entities than physical objects and their properties and relations; everything natural can ultimately be viewed in just those terms; and hence it is the methodological and explanatory canon of physics that defines the scientific enterprise as a whole.

Liberal naturalism, on the other hand, embodies a specific kind of ontological pluralism. It does not recognise any supernatural entities, since it abides with the core commitment of a naturalistic stance. But it does allow natural entities to come in different kinds which are not reducible one to the other. So liberal naturalism subscribes to ontological pluralism within the natural world: it recognises different kinds of natural entities as entities in their own right. With that comes the possibility of methodological and explanatory pluralism: ontological pluralism can not rule out, but it does not imply it either of course, that the study of different kinds of natural entities calls for different methodologies, and also that these different kinds of entities need to be explained in different ways.

---

[18] 'Conservative naturalism' comes close to what is also called 'classical' or 'scientific' naturalism; cf., De Caro & MacArthur (2004).

It will be clear what determines the choice between conservative and liberal naturalism: Are there natural phenomena that require for their proper explanation reference to entities that need physical realisations, so are not supernatural, but that can not be exhaustively characterised in terms of their physical realisations and their properties? If no, than conservative naturalism is a viable stance; if yes, then liberal naturalism is vindicated.

In the next section we will give a number of examples from a variety of disciplines that suggest that, yes, such entities do exists, and that hence, no, conservative naturalism is not a viable option in those disciplines. Moreover, we claim that linguistics is one of these disciplines, so after this brief review we will go into the question what this means for linguistics as a discipline.[19]

9 Groups on active duty: some examples

The first two examples come from biology. The first concerns the role that groups might play as vehicles of selection in the evolutionary process. Where the standard conception of evolution views individuals as the sole instruments in the selection process, some have argued that we need also to take into account the groups to which individuals belong, since the effects of behaviour on groups may be significant, and hence groups themselves may act as selection vehicles as well. Although the idea of group selection was entertained by Darwin, it had fallen in disrepute for a long time. It was revived by Wilson and Sober some twenty years ago, and since then it has been the subject of a lively debate. This is how Wilson and Sober motivate the idea in one of their earliest papers (Wilson & Sober 1994, p. 605–606):

> There are compelling intellectual and practical reasons to distinguish between be-
> haviours that succeed by contributing to group-level organisation and behaviours
> that succeed by disrupting group-level organisation. [...] A concern for within-
> group versus between-group processes characterises the human mind and should
> characterise the study of the human mind as well.

Since its original formulation, the theory of group selection has gone through some revisions and the debate about has not subsided. But the point for our discussion is

---

[19] To avoid misunderstanding, the examples from other disciplines serve to illustrative the viability of liberal naturalism as such, not as (additional) arguments for a liberal naturalistic stance in linguistics. Whether the latter is plausible or not does not depend on the former, of course.

that if this is on the right track, then selection is a 'multi-level' process. It involves not just individuals, but also groups, which hence have to be regarded as basic, non-reducible entities: physically realised in their constitutive individuals, but nevertheless with properties of their own.

The second example from biology concerns trait inheritance. Again, the standard theory is individualistic, as it locates the inheritance mechanisms solely in the individual genome, and thus considers that to be the only relevant factor for phenotypic expression. However, this view arguably ignores the obvious and important role of other contributing factors. This is how Dupré characterises the situation (Dupré 2014, p. 81–82):

> To expect in general that identifiable bits of the genome will have privileged relations to particular traits of the phenotype, given that they do not typically even have unique relations to particular functional proteins, would be hopelessly unrealistic.The notion of the genome as composed of a series of genes 'for' particular phenotypic traits has gone the way of phlogiston. […] The classificatory division of the genome within genomics proper, therefore, is one driven very much by theoretical considerations, and is little effected by social factors in the interesting sense of 'social'. If genomics eventually gives us a good understanding of development, then we might expect to derive real abilities to control developmental outcomes, human and otherwise. But given the demonstrable complexity of development and of its joint dependence on internal and environmental factors, the task is a daunting one.

This goes against a widespread belief that phenotypic expression of inherited traits is a matter of the individual only, whereas it is quite obvious that environmental factors of various kinds also play a key role.[20] And it seems quite likely that among these environmental factors are entities that are not individual in nature, such as kinship and other social groups, economic and social institutions, and so on.

Other scientific disciplines also provide arguments against the individualistic bias that is a characteristic feature of conservative naturalism. For example, in

---

[20] On the widespread nature of the individualistic bias, cf., Lobo & Shaw (2008):

> In an age in which scientists and the public are excited about the sequencing of the entire human genome, we need to temper that excitement, at least a little, and be careful not to believe all the hype surrounding genes' involvement in determining development and behaviour. While an organism's genetic makeup plays a critical role in its development, there is also a rich and complex interplay between the genome and cues from the environment.

economics and decision theory the importance of the role of groups, in particular of such factors as group knowledge and rationality, group intentionality and action, is increasingly acknowledged. And more and more it becomes clear that the group properties mentioned can not always be reduced to their individual counterparts. Witness the 'Diversity Prediction Theorem', a mathematical result concerning the accuracy of group and individual prediction:[21]

> The squared error of the collective prediction equals the average squared error minus the predictive diversity

What this says is that the prediction of the group as a whole becomes better as the diversity of the predictions of its members increases. Again, what we have here is a property of an entity that can not be reduced to properties of its constitutive members, which means that groups as entities need to be taken on board. And unlike its conservative counterpart, liberal naturalism allows for that.

And then of course there is an entire tradition in the social sciences that is naturalistic in its basic stance, yet embraces a social ontology in which such entities as communities, institutions, practices, and the like play an essential role. That tradition has many different faces, as the works of authors such as Foucault, Bourdieu, Giddens, Schatzki, and many others illustrates. But they all subscribe to the starting point that was formulated already in the nineteenth century by one of the founders of modern sociology, Emile Durkheim (Durkheim 1895):

> Here, then, is a category of facts which present very special characteristics: they consist of manners of acting, thinking and feeling external to the individual, which are invested with a coercive power by virtue of which they exercise control over him. Consequently, since they consist of representations and actions, they cannot be confused with organic phenomena, nor with psychical phenomena, which have no existence save in and through the individual consciousness. Thus they constitute a new species and to them must be exclusively assigned the term social. It is appropriate, since it is clear that, not having the individual as their substratum, they can have none other than society, either political society in its entirety or one of the partial groups that it includes – religious denominations, political and literary schools, occupational corporations, etc.

We take it that these examples, coming from a range of different fields, though not always uncontroversial within the respective disciplines, do illustrate the point

---

[21] Cf., Page (2007).

that the strict conservative naturalism that many would subscribe to as a matter of course is not the only option, and in some cases arguably not the right one. Liberal naturalism seems the way to go.

10   And in linguistics, too

In this section we list just a few of the many phenomena that substantiate the claim that what was claimed in the previous section holds for linguistics as well. The point here is not to introduce anything new, as these phenomena are well-known and extensively studied. Rather it is to emphasise that social entities, used here as a catch-all phrase to refer to groups, communities, social practices, institutions, and sundry entities, play a key role in an account of them. And that means that if we are to explicate the ontological implications of such accounts in a naturalistic setting we need a liberal naturalism.

First of all, there is the area of speech act theory from which reference to institutions and institutional practices can not be eliminated. Despite the fact that in most classical formulations the focus in the analysis of linguistic actions, and of the intentionality that is involved in them, remains at the individual level, it is also true that the institutional level is a necessary ingredient of a proper understanding of many speech acts. A similar argument can be made for the analysis of linguistic performativity, which also requires the acknowledgement of institutional frameworks. The ensuing ontological diversity is usually not spelled out very explicitly, but it is there.

A second example comes from semantics, and is provided by various forms of semantic externalism and by the phenomenon of the division of linguistic labour. By its very presuppositions, semantic externalism is at odds with the cognitivist conception of meaning that locates meaning in the individual's mind (or brain, as the case may be). It brings in other factors that co-determine meaning, such as the physical world in the case of natural kind terms, or the social or cultural environment in other cases. Especially the latter are hard to account for on a strictly individual basis. And the same holds for the phenomenon of division of linguistic labour, with it associated references to different subgroups in a linguistic community. The commitment to a social ontology and the concomitant failure of a conservative naturalistic stance may not always be very clear from the formal accounts that have been proposed, but they are there.

At the interface between semantics and pragmatics we find a third phenomenon: that of meaning contextualism. The assumption of stable ('literal') meanings is a prerequisite for the classical individualistic sender – receiver model of linguistic communication to make sense. It does provided that meanings can be regarded as individual assets. For only on that assumption does it make sense that meanings are used (expressed, transmitted) in context. But if meanings themselves depend on context, such an account is no longer an option. Meaning contextualism reverses the order of explanation: it is not individual, stable meanings that explain communication, rather it is the other way around, it is communicative practices that produce meanings that have (limited) stability. This is how Medina puts it (Medina 2004, p. 571):

> [S]emantic determinacy is the always fragile and relative accomplishment of communicative interactions which rest on a tacit agreement in action that is always undergoing transformation. Meanings become contextually determinate through the practical consensus achieved by participants in situated linguistic interactions against the background of shared practices.

And again, it is only when we spell out meaning contextualism that we find that we can not make do with a strictly individualistic model, but that we need to accept communicative practices and the communities in which they exists as bona fide denizens of our theoretical framework. The point can be strengthened further by pointing to such a phenomenon as successful communication with incomplete understanding. Just like the phenomena just mentioned this is something that lies outside the grasp of a linguistic theory that embraces conservative naturalism and its strict individualism.[22]

---

[22] At this point it may be good to point out that in many explicitly non-Chomskyean paradigms, in which semantics and pragmatics are considered to be intrinsic parts of linguistic theory, methodological individualism still plays a role. Here is an example from construction grammar, cf., the following passage in Goldberg (1995):

> Theorists working in this theory [i.e., construction grammar, MS-MVL] share an interest in characterising the *entire* class of structures that make up language, not only the structures that are defined to be part of 'core grammar'. This interest stems from the belief that fundamental insights can be gained from considering such non-core cases, in that the theoretical machinery that accounts for non-core cases can be used to account for core cases. [...] Construction Grammarians also share an interest in accounting for the conditions under which a given construction can be used felicitously, since this is take to be part of the speakers' competence or knowledge of language; from this interest stems the conviction that subtle semantic and pragmatic factors are crucial to understanding the constraints on grammatical constructions.

11 Liberal naturalism and the dispositional view

So it seems that a comprehensive account of language needs the liberal natural-istic perspective, i.e., a framework that can accommodate both individual and social dimensions of language and language use, and that is not committed to the hypthetico-structuralist model as the only respectable type of explanation, and that thus makes room for the incorporation of dispositional properties at the individual level and for social entities as autonomous elements in its ontology.

Of course groups, social practices, and so on, are not not supernatural entities, they have physical realisations in the natural world. What is important to note here is that, in general, the physical realisations of these social entities are not ontologically homogeneous. In some cases, specific parts of such a realisation may be identified with individual language users. But others pertain to quite different kinds of entities, such as implements, historical records, spatial configurations, and other material aspects. And yet others need to be explained in terms of properties of collectives, such as group intentionality and group knowledge.

This has important consequences, since it paves the way for ontological plu-ralism, viz., the acknowledgement of different categories of natural entities that are irreducible one to the other. Reduction of behavioural and psychological properties of groups and institutions and their practices, to properties of the individuals that realise them is not always possible. And even the individual behavioural and psychological properties that are involved cannot be reduced, as conservative natu-ralism would require, to underlying structural properties at the neurophysiological level, at least not in a significant sense.

Thus the picture that emerges is pluralistic in a number of ways. First of all, it comprises different kinds of individual entities, not just human language users. Second, it acknowledges collectives of such individuals, such as linguistic communities, as entities in their own right. And thirdly, it allows not just for the structural properties of the entities involved to play an explanatory role, but also for their dispositions.

This requires an explanatory model that can integrate these various aspects in such a way that the different components and their interactions can all be

---

Obviously, the scope of linguistic theory as it is conceived of here is wider than in the generative paradigm. But note that the central role of the individual is still present: linguistic competence may be a more encompassing notion but it still remains an individual affair

accounted for. An example of such a model, taken from Vanderbeeken & Weber (2002), is given in figure 1.
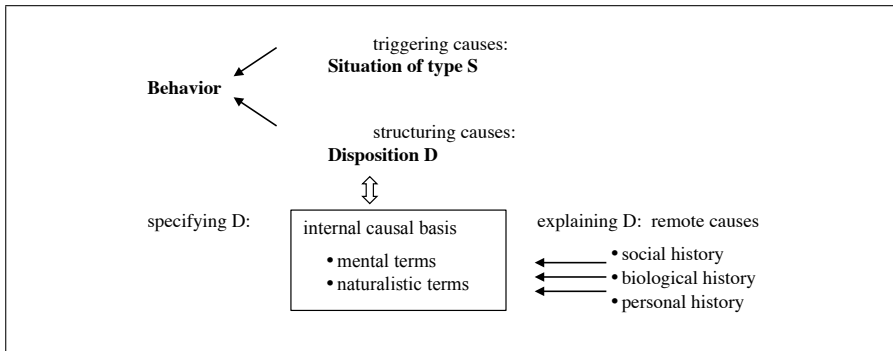

Figure 1: Dispositions in explanations

Not only does a model such as this make room for dispositions as an explanatory factor, it also allows for an account of the complex interactions between different kinds of causal factors, among which are dispositions, structural factors, and environmental factors. A model along these lines thus seems much more adequate for dealing with the rich and varied phenomena that language and language use are.

12   The role of linguistic theory in liberal naturalism

But what, one may (and should) ask, is the role of linguistic theory in such a complex model?

Marr (Marr 1977) usefully distinguished between two types of theories of cognitive phenomena. 'Type 1' theories are theories in which one can distinguish Marr's famous three levels. The top level is like the competence models of linguistics – it specifies in mathematical terms the inputs and outputs of a cognitive process (such as, on one view, language production and comprehension) viewed as an information processing task. Once the top level is specified one can construct algorithms meeting the specifications and compare these algorithms with respect to neural implementability. It is characteristic of a type 1 theory that its inputs and outputs can be described formally, and this makes it doubtful whether linguistics can be regarded as a type 1 theory. Before we substantiate these doubts, we quote

114

Marr's very interesting description of the alternative, 'type 2' theories, which apply when (Marr 1977, pp. 38–39):

> [. . .] a problem is solved by the simultaneous action of a considerable number of processes, whose interaction is its own simplest description.[. . .] At each moment only a few of the possible interactions will be important, but the importance of those few is decisive. Attempts to construct a simplified theory must ignore some interactions; but if most interactions are crucial at some stage [. . .] a simplified theory will prove inadequate.

Indeed, in language comprehension, the input may itself be the product of interacting processes, in the following sense: a piece of discourse may seem uninterpretable gibberish (even though not obviously ungrammatical), until interaction with a non-verbal information source restores sense to the discourse. A classic example of this phenomenon is the experiment in Bransford & Johnson (1972), in which subjects were presented with the following discourse, supposedly from a man muttering to himself:

> If the balloons popped, the sound wouldn't be able to carry since everything would be too far away from the correct floor. A closed window would also prevent the sound from carrying, since most buildings tend to be well insulated. Since the whole operation depends on a steady flow of electricity, a break in the middle of the wire would also cause problems. Of course, the fellow could shout, but the human voice is not loud enough to carry that far. An additional problem is that a string could break on the instrument. Then there could be no accompaniment to the message. It is clear that the best situation would involve less distance. Then there would be fewer potential problems. With face to face contact, the least number of things could go wrong.

Some subjects were shown figure 2 before they had to interpret the discourse.
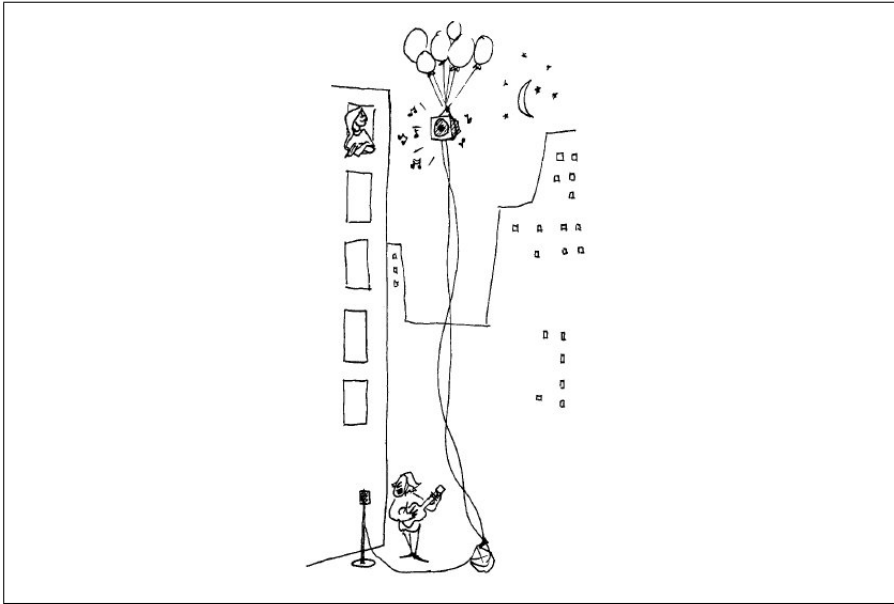
115

Figure 2: A modern serenade

These subjects could adequately recall and summarise the discourse. By contrast, subjects who were shown the picture after the discourse or not all, were unable to summarise or even recall the discourse. Thus the hearer must have access to the speaker's 'situation model' to be able to interpret the discourse. Realistically, this means that the hearer, starting from a few clues concerning the speaker's situation model, elaborates the situation model and interprets the discourse in an interactive and defeasible manner. Since making the picture available after the speaker has stopped talking does not benefit the hearer, processing must obey the 'principle of immediacy' , i.e., 'all available information will immediately be used to co-determine the interpretation of the speaker's message' (Hagoort & Berkum 2007, p. 801). If one adopts this processing principle, linguistics cannot be a type 1 theory.

13   Conclusion

Although much more detail needs to be added, we do think that the considerations given in this paper support the conclusion that an adequate account of language and language use needs a liberal naturalistic perspective. The main reason is

that language manifests itself in categorically distinct ontological entities, and methodological individualism does not succeed in constructing an ontologically homogeneous base for this variety. We have to accept that language is ontologically heterogenous and explanatorily not uniform. And the way in which a conservative naturalistic view is defended in the literature supports the claim that theories that aim to give a uniform explanation based on a homogeneous ontology are informed by some form of ideologically informed reductionism.

An alternative, liberal naturalistic view, which is ontologically, methodological and explanatory heterogeneous seems to be called for, and a dispositional account of several aspects of language and language use needs to be incorporated in such a view. The role of linguistic theory in this alternative view is to provide a systematic description of certain properties of linguistic expressions that can be part of such an encompassing account of language and its use. This means that linguistic theory must be viewed as one methodology among many, that deals with one particular aspect among many, of a heterogeneous phenomenon. Viewed in this way linguistic theory offers a conceptual reconstruction of certain aspects of language that also may suggest further empirical investigations. That is an important element of our investigations, but not the final word: linguistic theory is not an explanatory theory of a distinct empirical phenomenon.

References

Bigaj, Tomasz. 2012. Ungrounded dispositions in quantum mechanics. *Foundations of Science* 17. 205–21.

Bransford, J.D. & M.K. Johnson. 1972. Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior* 11. 717–26.

Chomsky, Noam. 1995. Language and nature. *Mind* 104(413). 1–61.

Chomsky, Noam. 2000. *New horizons in the study of language and mind.* Cambridge: Cambridge University Press.

Culicover, Peter W. & Ray Jackendoff. 2005. *Simpler syntax.* Oxford: Oxford University Press.

De Caro, Mario & David MacArthur. 2004. Introduction: The nature of naturalism. In Mario De Caro & David Macarthur (eds.). *Naturalism in question.* 1–17. Cambridge, MA: Harvard University Press.

Dupré, John. 2014. What's the fuss about social constructivism? *Episteme* 1(1). 73–85.

Durkheim, Émile. 1895. *Les règles de la méthode sociologique.* Paris: Librairie Felix Alcan.

Field, Hartry. 1980. *Science without numbers: A defense of nominalism.* Princeton, NY: Princeton University Press.

Givon, Talmy. 2013. The intellectual roots of functional linguistics. In Shannon Bischoff & Carmen Jany (eds.). *Functional approaches to language.* 9–30. Berlin: De Gruyter Mouton.

Goldberg, Adele. 1995. *A construction grammar approach to argument structure.* Chicago, Ill.: The University of Chicago Press.

Hagoort, Peter & Jos van Berkum. 2007. Beyond the sentence given. *Philosophical Transactions of the Royal Society B* 362. 801–11.

Hinzen, Wolfram & Juan Uriagereka. 2006. On the metaphysics of linguistics. *Erkenntnis* 65. 71–96.

Lobo, Ingrid & Kenna Shaw. 2008. Phenotypic range of gene expression: Environmental influence. *Nature Education* 1(1). 12.

Marr, David. 1977. Artificial intelligence: A personal view. *Artificial Intelligence* 9. 37–48.

McMullin, Ernan. 1978. Structural explanation. *American Philosophical Quarterly* 15(2). 139–47.

Medina, José. 2004. Anthropologism, naturalism, and the pragmatic study of language. *Journal of Pragmatics* 36(3). 549–73.

Page, Scott E. 2007. *The difference: How the power of diversity creates better groups, firms, schools, and societies.* Princeton: Princeton University Press.

Stenning, Keith & Michiel van Lambalgen. 2008. *Human reasoning and cognitive science.* Cambridge, MA: MIT Press.

Stokhof, Martin & Michiel van Lambalgen. 2011a. Abstraction and idealisation: The construction of modern linguistics. *Theoretical Linguistics* 37(1–2). 1–26.

Stokhof, Martin & Michiel van Lambalgen. 2011b. Comments–to–comments. *Theoretical Linguistics* 37(1–2). 79–94.

Suarez, Mauricio. 2007. Quantum propensities. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 38(2). 418–38.

Thompson, Ian J. 1988. Real dispositions in the physical world. *The British Journal for the Philosophy of Science* 39. 67–79.

Vanderbeeken, Rob & Erik Weber. 2002. Dispositional explanations of behavior. *Behavior and Philosophy* 30. 43–59.

Wilson, David Sloane & Elliot Sober. 1994. Reintroducing group selection to the human behavioural sciences. *Behavioural and Brain Sciences* 17(4). 585–654.

Authors

Martin Stokhof
ILLC/Department of Philosophy
University of Amsterdam
m.j.b.stokhof@uva.nl

Michiel van Lambalgen
ILLC/Department of Philosophy
University of Amsterdam
m.vanlambalgen@uva.nl