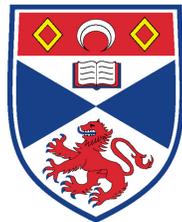


The Foundations of Linguistics: Mathematics, Models, and Structures

Ryan M. NEFDT

*This thesis is submitted in partial fulfilment for the degree of
Doctor of Philosophy
2016*

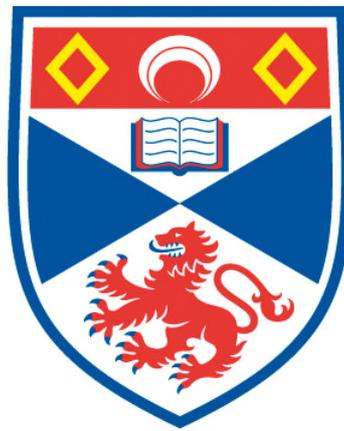


University
of
St Andrews

**THE FOUNDATIONS OF LINGUISTICS:
MATHEMATICS, MODELS, AND STRUCTURES**

Ryan M. Nefdt

**A Thesis Submitted for the Degree of PhD
at the
University of St Andrews**



2016

**Full metadata for this item is available in
St Andrews Research Repository
at:**

<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:

<http://hdl.handle.net/10023/9584>

This item is protected by original copyright

**This item is licensed under a
Creative Commons Licence**

Declaration

I, Ryan Mark Nefdt, hereby certify that this thesis, which is approximately 78 806 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in January 2014 and as a candidate for the degree of Doctor of Philosophy in January 2014; the higher study for which this is a record was carried out in the University of St Andrews between 2014 and 2016.

Candidate's signature

31/08/2016

Date

I, Ephraim Glick, hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Supervisor's signature

Date

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

Candidate's Signature

Supervisors' signatures

This page intentionally left blank

Contents

Abstract	vii
Acknowledgements	viii
Preface	xii
I Three Grades of Mathematical Involvement in Linguistics	1
1 The <i>Mathematisation</i> of Natural Language	2
1.1 Formalisation, Mathematisation and Conservativeness	3
1.2 Hilbert, Bar-Hillel and Chomsky	4
1.3 The Legacy of <i>Syntactic Structures</i>	8
2 Methodology and Ontology	14
2.1 Linguistic Conceptualism	15
2.2 Linguistic Platonism	20
2.3 Linguistic Pluralism	25
3 Three Grades of Mathematical Involvement	30
3.1 The First Grade: Models and Linguistic Reality	35
3.2 The Second Grade: Representational Realism	48
3.3 The Third Grade: Grammars as Mathematical Theories	57
3.4 Subconclusion	61
II Methodology in Syntax and Semantics	62
4 Modelling in Syntax	63
4.1 Modelling and Idealisation	64
4.2 How the Laws of Linguistics Might Lie	65
4.3 Minimalist Idealisation from ST to Minimalism	66
4.3.1 Minimal Generation	68

4.3.2	Isolation	72
5	The Dynamic turn and other Frameworks	77
5.1	Other Generative Frameworks	77
5.2	Modelling Dynamics	80
5.3	Model-theoretic Syntax and Overgeneralisation	83
5.4	Subconclusion	87
6	Modelling in Semantics	89
6.1	Formal Semantics as Applied Mathematics	89
6.2	Aims, Target and Methods	91
6.2.1	Semantic Explananda	91
6.2.2	Target of the Model	95
6.2.3	Methods and Model Theory	101
6.3	On the Unreasonable Effectiveness of Semantics	105
6.4	Predicting Syntactic Felicity in the Real World	111
6.5	The Syntax-Semantics Interface	115
6.5.1	Semantics without Semantic Value	120
6.5.2	Direct Modelling	123
6.6	Subconclusion	125
7	Part Conclusion: Structural Realism	126
III	Languages and Other Abstract Structures	131
8	Three Desiderata	134
8.1	Linguistic Creativity and Infinity	134
8.2	Of Tokens and Types	140
8.3	Realism and Respect	143
8.3.1	Taking Stock	148
9	Against Platonism	150
9.1	The Right Kind of ‘Wrong View’	151
9.2	Benacerraf’s Dilemma and Respect	154
9.3	Conceptual Distinctness	156
10	<i>Ante Rem</i> Realism	159
10.1	Mathematical Structuralism	159
10.2	Linguistic Structures	164
10.3	Quasi-Concreteness	165
10.4	Banishing a Dogma	171
10.4.1	Types as Archetypes	174

10.4.2 Types as Wholes	178
10.4.3 Types and the Representation Relation	182
10.5 Subconclusion	189
Conclusion	191
References	194

Abstract

The philosophy of linguistics is a rich philosophical domain which encompasses various disciplines. One of the aims of this thesis is to unite theoretical linguistics, the philosophy of language, the philosophy of science (particularly mathematics and modelling) and the ontology of language. Each part of the research presented here targets separate but related goals with the unified aim of bringing greater clarity to the foundations of linguistics from a philosophical perspective.

Part I is devoted to the methodology of linguistics in terms of scientific modelling. I argue against both the Conceptualist and Platonist (as well as Pluralist) interpretations of linguistic theory by means of three grades of mathematical involvement for linguistic grammars. Part II explores the specific models of syntactic and semantics by an analogy with the harder sciences. In Part III, I develop a novel account of linguistic ontology and in the process comment on the type-token distinction, the role and connection with mathematics and the nature of linguistic objects.

In this research, I offer a structural realist interpretation of linguistic methodology with a nuanced structuralist picture for its ontology. This proposal is informed by historical and current work in theoretical linguistics as well as philosophical views on ontology, scientific modelling and mathematics.

Acknowledgements

I am greatly indebted to Ephraim Glick and Josh Dever for their dedicated and excellent supervision throughout my time at the University of St Andrews, the Arché Research Centre and the University of Texas at Austin. Ephraim looked at countless drafts of various parts of this thesis when they were still in their infancy. His comments and criticism were always of an extremely high standard. He has contributed to my development as a philosopher in so many significant ways. His guidance also often extended beyond the theoretical as he provided me with invaluable assistance on practical matters pertaining to the job market and general academic life. Josh was essential to me seeing the bigger picture but also working through the details of specific examples with his characteristic mixture of precision and insight. I also appreciate all the time he made for me at UT Austin.

At Arché, I am grateful for the help and guidance of Arché faculty, fellow researchers and visitors: Derek Ball, Patrick Greenough, Stephen Read, Andy Egan, Jody Azzouni, Mark Thakkar, Ravi Thakral, Poppy Mankowitz, Claire Field, Matthew McKeever, Matthew Cameron, Asgeir Matthíasson, Bruno Jancinto, Fenner Tanswell and others.

I would also like to heartily thank fellow participants of the Foundations of Linguistics workshop at the Braunschweig Institute of Technology in 2015. Geoff Pullum, David Pitt, Terry Langendoen, Scott Soames, Andras Kornai, Bob Levine, Christina Behme, Hans Heinrich Lieb and Martin Neef were all helpful in commenting on or reading work that would eventually find its way into this research. I thank David Pitt for additionally going over some of the details of the ontology of natural language with me at Cal. State, LA and enlightening me about the views and opinions of the late Jerrold Katz. I am also especially grateful to Geoff Pullum, at Edinburgh's esteemed linguistics department, for his detailed comments on a draft of this manuscript and taking the time to meet with me on several occasions to discuss various aspects of the history of formal syntax and formal language theory generally. He remains a personal hero of mine.

During my time at Yale University in 2016, Zoltan Szabó's advice and support proved excellent, especially for the chapters on semantic modelling and realism. I would also like to thank Kate Stanton for her inimitable assistance on the content of Part I. Lastly at Yale, I would like to thank participants of the Semantics Sem-

inar in the philosophy and linguistics departments for their helpful comments on a presentation of my research there.

Geoff Poole at the Centre for Research in Linguistics and Language Sciences at the University of Newcastle deserves my thanks for his assistance with some aspects of my analysis of generative syntax, Bernhard Weiss at the Department of Philosophy at the University of Cape Town for his insightful tutelage during my MA degree in philosophy and continued support thereafter, and Reinhard Blutner and Henk Zeevat both of the Institute for Logic, Language and Computation at the University of Amsterdam, where I conducted my MSc in Logic and Linguistics, for the requisite formal training in formal syntax, semantics and optimality theory. Reinhard also guided some of my thoughts on Part II.4. In addition, Jarda Peregrin of the Czech Academy of Sciences was a much needed voice of caution and insight on the connection between structuralism and semantics (particularly of an inferentialist persuasion).

At the philosophy department at the University of Michigan, I would like to thank Lawrence Sklar for introducing me to various issues in the philosophy of science and physics and Sarah Moss for my first contact with formal methods (outside of logic) in philosophy and semantics during my year as a research scholar there and subsequent period as a visiting graduate student. I also thank Hans Kamp, at UT Austin, for shedding necessary light on some issues related to early formal semantics and some of the scientific motivations behind Discourse Representation Theory. Finally, to my good friend Tanmay Inamdar at the department of Mathematics at the University of East Anglia, a special thanks for always being receptive to the questions and worries of a philosopher concerning your fields of set theory and mathematical logic respectively.

I would like to sincerely thank audiences in St Andrews, Stirling, Edinburgh, Leeds, Newcastle, Lodz, Szklarska Poreba, Amsterdam, Braunschweig, Boston and New Haven in both linguistics and philosophy departments for their useful comments and suggestions during my conference and research travels.

Lastly, I would like to thank my sister Taryn, the scientist, for leading the way and my partner, Anastasia, for her fantastic support during the difficult periods of this research.

To my parents, Anthea and Gordon, for their eternal support and encouragement without which none of this wonderful journey would have been possible

“Generative linguistics is a perplexing field of inquiry when looked at with philosophic issues in mind.”

– Sylvain Bromberger, *Reflections on Chomsky* (1989)

“On the structuralist perspective, what is an object in a given context depends on what concepts or predicates are in use, and this depends on what concepts or predicates are available. It is through language that we organize the world and divide it into objects.”

– Stewart Shapiro, *Philosophy of Mathematics: Structure and Ontology* (1997)

“Thus in the scientists’ use, ‘model’ denotes what I would call a model-type. Whenever certain parameters are left unspecified in the description of a structure, it would be more accurate to say (contrary of course to common usage and convenience) that we described a structure-type.”

– Bas van Fraassen, *The Scientific Image* (1980)

Preface

The philosophy of linguistics is a rich philosophical domain which encompasses aspects of the philosophy of language, theoretical linguistics, cognitive science, science and mathematics. Traditionally, the study of the methodological and ontological consequences of the field has been conducted within the remit of theoretical linguistics. The “linguistic approach” has led to many misconceptions and confusion within the foundations of linguistics debate. For instance, I argue that a symptom of the linguistic approach is the vast *mathematisation* which has infected the field to such an extent that Platonist theories are only the logical conclusion of comments and claims made in the standard (mentalist) picture of the aim of the science.

In this thesis, I will work within a more “philosophical approach” to certain questions in the foundations and the ontology of linguistics. This approach draws from the fields of the philosophy of science (in Part II), the philosophy of mathematics (in Part III), metaphysics and the philosophy of language to engage with central questions in theoretical linguistics in a philosophically rigorous manner. Despite a marked departure from the traditional approach, I do maintain a close connection with actual (first-order) linguistic research throughout so as not to diverge from the goal of explaining the science as it is and its objects as they are, i.e. the interpretive project.¹

This thesis takes on a tripartite structure. The first theme which will receive treatment throughout will be related to how mathematics has played a role in shaping contemporary linguistics and how it should be interpreted within linguistics. The next theme is related and aims to establish a scientific modelling approach to the understanding of linguistic grammars. I will not only argue that this approach is appropriate but also that it solves many puzzles related to the concept of grammar prevalent in the syntax and semantics literature. Instead of resulting in an anti-realist position on the scientific status of linguistics, however, I will suggest a structural realist foundation. Parts I and II of the thesis can be seen as an attempt to explain the nature of linguistic models *qua* structures, while part III precisifies the

¹This approach is by no means unique to this work but rather a rich tradition I hope to follow in. Pullum, Scholz, George, Devitt, Pelletier, Peregrin, Szabó, Partee, Soames, Stainton, Tomalin and many others are excellent examples of philosophical rigour coupled with an emphasis on serious interpretation of linguistic theory.

realist component of the view in a novel way. Thus, the overall plan of the thesis is to provide novel insights into the structures involved in linguistic theorising and modelling and in so doing shed light on their methodological, ontological, scientific and mathematical significance.

One of the major claims of this thesis is that linguistic methodology and its ontology are separable and should be treated independently. Part I will be devoted to establishing this claim. I argue that if this point is not appreciated, many confusions ensue. I do so by suggesting three grades of mathematical involvement in linguistics in terms of the notion of a grammar. I advocate the first grade which views grammars as ontologically neutral and akin to scientific models in their methodological status. In chapter 2 of Part I, I will proffer this novel interpretation of theoretical linguistics in accordance with scientific modelling. In Part II, I will develop a thorough account of linguistic modelling practices in both formal syntax and semantics in light of the aforementioned interpretation of the core notion of a grammar. This account will show continuity with modelling practices in the other sciences such as physics, biology and economics.

Once my methodological stance has been argued, I will devote my attention to the ontology of linguistics and natural language. I offer a Realist approach to linguistic ontology, one which draws from the Platonist ontology of Katz, Postal and Soames but diverges significantly in terms of their metaphysics. Specifically, I will argue for a mathematical structuralist account similar to the theories of Shapiro (1997) and Resnik (1997) for mathematical reality. My picture departs from theirs in important ways though. For instance, I eschew the abstract ontology of structures in favour of quasi-concrete structures which have parts in physical reality (Parsons 1990). I further amend this picture to accommodate a more naturalistic type-token distinction in line with Szabó (1999). Thus, linguistics on this view is seen as a quasi-empirical discipline which deals with phenomenon causally related to both mental competence in language (I-languages à la Chomsky) and more formal aspects of its description.

It is important to note that although much of the discussion of linguistics, especially syntax, will draw from the work of leading linguists such as Chomsky, I do not plan to engage in any serious Chomsky exegesis directly. My plan is rather to assess and resolve certain philosophical puzzles presented by the alleged scientific nature of the field and the mathematics which underlies it as put forward by its practitioners.

Parts of this thesis can be considered to be exercises in the philosophy of science (certainly Part II). Parts are concerned more with ontological issues and the philosophy of mathematics (Part III) and yet others are combinations of the aforementioned fields and the philosophy of language and linguistics (Part I). Throughout, the aim is not only to shed new light on old debates but also to offer genuinely novel accounts of linguistics as a science and natural language as a object or rather structure in the natural world.

This page is intentionally left blank

Part I

Three Grades of Mathematical Involvement in Linguistics

Chapter 1

The *Mathematisation* of Natural Language

The ontological basis of the linguistic enterprise has been contested since something resembling an official stance was adopted through the generative or biolinguistic tradition in the late 1950's. On this view, natural languages are states of the mind/brain and thus part of cognitive-psychological reality. An alternative approach (proffered by Katz 1981, Langendoen and Postal 1984, Soames 1984, Katz and Postal 1991, and Postal 2003) considers natural languages (and the sentences of which they are comprised) to be abstract objects, in the sense of being mind-independent and non-spatio-temporally extended. Linguistics on this view is a formal science on the level of logic or mathematics. Arguments have gone back and forth in favour of specific views without much common ground, to the effect that these linguists and philosophers often seem to talk past each other. In this chapter, my moderate aim will be to reformulate the debate such that a clear battleground is demarcated (at least). My chief aim will be to offer a non-ontologically-committing view which draws on insights from computational linguistics and the philosophy of science. In order to (re)situate the debate, I shall first identify the problem of 'mathematisation' in section 1 and then describe the main positions on the foundations of linguistics with relation to it in section 2, and in section 3 suggest three grades of mathematical involvement for the grammars of linguistic theory. These grades involve the methodological attitudes linguists take towards the grammar. More specifically, the grades correspond to possible positions on the nature of the mathematical apparatus used in the grammar and its relation to the nature of natural language itself. Each grade has its own set of difficulties but I suggest that the first grade of involvement offers us the most neutral and reasonable approach to linguistic methodology (while avoiding certain issues within its ontology). The last two grades are argued to be more philosophically fraught and thus unstable foundations for the field (at least in their current forms).

1.1 Formalisation, Mathematisation and Conservativeness

I begin this section with a distinction. The distinction is between the concepts of *formalisation* and *mathematisation* respectively. Formalisation is the familiar tool of simplifying natural structures or phenomena for the purpose of making them more amenable to precise characterisation, often in terms of the language of first-order logic. This tool is certainly wide-spread in linguistics and philosophy. Mathematisation, on the other hand, can be seen as the process of rendering natural phenomena into mathematical structures or entities by either formal modelling, analogy or alleged proof. It marks an ontological shift in the target system (or sometimes omission of the target altogether). In other words, it approaches the subject of linguistics as a mathematical puzzle capable of precise mathematical characterisation and resolution without any attempt to interpret any features within the target system. In what follows I shall argue, for instance, that an example of the formalisation is the use of lambda abstraction as a means of variable binding and functional application in semantics. While an example of mathematisation are present in claims such as “the processes described by generative grammar are functions computed over physical symbols by a Turing machine implemented in the human brain” (Hinzen and Uriagereka, 2006: 71).

In this section, I want to provide some details concerning the concept of mathematisation, specifically within the context of the linguistic project. The concept will remain in the background of most of the discussions to follow. Unfortunately, the term has not received any proper definition, despite some scattered usage across reflective theoretical work in the philosophy of science. I will provide a working definition for present purposes but this characterisation should not be considered definitive but rather illustrative. For the sake of contrast, I will begin with the more familiar concept of formalisation.

The concept and technique of formalisation, its features and fecundity have been addressed in philosophy, logic and other fields. In general, despite cautionary tales, formalisation is considered benign and mostly useful. As Pullum (2013) states “formalization is the use of appropriate tools from mathematics and logic to enhance explicitness of theories [...] Any theoretical framework stands to benefit from having its content formalized” (493). One important feature of formalisation is conservativeness, a property which received a more controversial treatment in Field (1980) and the mathematical nominalism with which it came. Contrary to popular indispensability arguments as to the essential place of mathematics within the natural sciences, Field proposed a thoroughgoing fictionalism about mathematical entities related to the concept of conservativeness given below.

A mathematical theory S is *conservative* if, for any nominalistic assertion A and any body of such assertions N , A is not a consequence of $N+S$ unless A

is a consequence of N alone (Field, 1985: 240).

The basic idea is that nothing that can be said with mathematics can not be said without it within scientific investigation. Whether or not this principle holds or can be shown to hold for the relationship between mathematics and the sciences remains to be seen (see Shapiro 1983, 1997). However, formalisation is certainly conservative in the way discussed above. Formal languages are devices for the simplification of an intended target domain. As the term suggests, the technique is meant to home in on the “form” of a problem and highlight the relations in a non-obfuscating manner. In fact, formalisation is related to an emphasis on a syntactic analysis of mathematical and scientific discourse, initially proposed in the philosophical project that came along with Hilbert’s programme in the early 20th century.

1.2 Hilbert, Bar-Hillel and Chomsky

For Formalists, mathematics was not a pre-interpreted theory of some extra-physical or mental reality but rather an uninterpreted calculus of symbols, the manipulation of which yields structures capable of later interpretation.

Every science takes its starting point from a sufficiently coherent body of facts given. It takes the form, however, only by *organizing* this body of facts. This organization takes place through the *axiomatic method*, i.e. one constructs a *logical structure of concepts* so that the relationships between the concepts correspond to relationships between the facts to be organized (Hilbert, 1899 [2004]: 540).

Hilbert’s axiomatic method was based on a concept of implicit definition. For example, unlike the axiomatics of Euclid which involved explicit definition of geometric terms such as a point being defined as “extensionless” or the like, Hilbert’s axioms introduced implicit definition directed toward the goal of divorcing theory from intuition (although intuitions might still play a motivating role for the axioms). As Shapiro puts it “geometry was becoming less the science of space or space-time, and more the formal study of certain structures” (2005: 63). The idea is that lines and points are to be defined purely in terms of the axioms of geometry and furthermore anything that fulfills the conditions set by the axioms will do equally well (what Shapiro calls “free-standing”). Nothing logico-conceptual is given in advance of theory. Thus, early proof theory took the shape of providing consistency proofs for parts of elementary arithmetic and analysis *via* the goal of establishing the consistency of the axioms (i.e. showing that they do not lead to contradiction). What separated Hilbert’s programme from other versions of Formalism was a focus on finitary methods, an aspect which took on special significance for linguists in the early 20th century. A discussion of finitism or strict finitism would take us too far afield.

Suffice to say, the beginnings of proof theory in mathematics incorporated a notion of consistency paired with existence and a distinct axiomatic approach coupled with the implicit definition of core concepts. The full fruition of this project could not be achieved since the consistency of arithmetic could not be proven through finitary means (thanks to Gödel's incompleteness result).¹ Nevertheless, Hilbert's axiomatic method had a profound effect on the scientific community at the time. This effect did not go unfelt within the linguistic community at the time. Early 20th century linguistics saw the anthropological goals of figures such as Sapir and Whorf take on a secondary role to the rigour of mathematical methodology.

[D]uring the 1930s and 1940s other developments in the theory of logical syntax occurred, which were ultimately to have profound implications for linguistic research, and the starting point was usually Hilbert's proof theory, which seemed to imply the meaning-less syntactic manipulations could suffice to resolve a whole range of epistemological problems (Tomalin, 2007: 89).

Bloomfield was among the first to embrace this approach to linguistics. In the spirit of the clarification of confusion, the identification of errors and the general precisification of the field, Bloomfield proposed the axiomatic method as its chief tool of investigation. Unfortunately, logical positivism and its circumscribed philosophical agenda also crept into Bloomfieldian linguistics. Meanings and any mental characterisation was anathema to early linguists. Formal syntax, just as the propositional or predicate calculus, was a more secure footing upon which to base the scientific study of language. Thus, the mathematical foundations became firmly entrenched.

One of the first sights of the application of recursive techniques or proof-theory in modern linguistics was in a paper by Bar-Hillel in 1953. Bar-Hillel attempted to extend the use of formal recursive techniques beyond the purely mathematical. By taking English as the metalanguage and French as the object language, he recursively redefined the basic parts of speech (noun, verb, etc) in order to establish a mathematical two-part recursive definition of a "proper" or grammatical French sentence.

Bar-Hillel's use of recursive definitions to analyse the structure of sentences in natural language can be viewed as one manifestation of this pervasive desire for the mathematisation of syntactic analysis, which became such a characteristic feature of certain kinds of linguistic research in the mid-twentieth century (Tomalin, 2007: 67).

¹However, the incompleteness proof did not destroy proof theory itself. Gentzen developed a proof-theoretic approach without the limitations of the Hilbert programme and finitism in full view of Gödel's incompleteness. His more specific aim was to prove the consistency of logical deduction within arithmetic.

This result did not go unnoticed by early Chomsky.² With this methodology came a movement away from empirical discovery procedures (the likes of which his mentor Harris had sort) toward mathematical precision and specification of evaluation procedures. This move was not unlike the move attributed to Hilbert in his axiomatic treatment of geometry, based not on spatial considerations, but pure mathematical structure arrived at through implicit definition. In the same way, geometry and linguistics can be thought to be motivated by intuition, but their study is ultimately tied up with mathematical investigation.³ Continuing with the theme of mathematisation, Peregrin (1995) describes Chomsky's contribution in the following way.

Chomsky's *novum* was that he proposed organizing the rules into a hierarchical system allowing the systematical generation, and basing all this upon setting up of the grammar as a real mathematical structure. Such a mathematization entailed an exceptional increase of rigour and perspicuity and, moreover, it led to the development of a metatheory, investigating into the formal properties of grammars (e.g. their relative strengths) (88).

Although, Hilbert, Bloomfield and Bar-Hillel all had an influence on the mathematical trajectory of the field. Chomsky's "novum" was more directly inspired by

²Tomalin traces a quasi-empiricist position of early Chomsky, inspired by Goodman's constructive nominalism and constructional system theory, to the eventual rejection of discovery procedures in favour of the recursive proof theoretic techniques of generative grammar.

³I will return to this analogy in section 3.3. On the philosophical side, I think that mathematisation is related to the treatment of the rules posited initially for the sake of modelling a phenomenon and then eventually for their own sake. A similar albeit Wittgensteinian diagnosis of the issue within the context of the foundations of linguistics can be found in Wright (1989). Simply put, it is a problem of rule-following. While generative grammar might have started out with a rather benign use of formal techniques in terms of formalisation and modelling. From the use of recursive structures and rules as tools, the movement placed the features of recursion and eventually merge at the forefront of its scientific agenda suggesting a strong relationship with discrete mathematics and especially arithmetic. With Platonism, mathematisation reached full fruition and the lines between mathematical objects and linguistic ones were all but dissolved. The problem of mathematisation can essentially be viewed as the progression and instantiation of the myth of the autonomy of rules or "the image of a rule as a rail laid to infinity" (Wright, 1989: 238). By placing rules or a generative grammar at a deep level of cognitive embedding or stating that natural language cognisers have "internalised" rule-systems, generative linguists commit this error which ultimately leads to Platonism. Consider the following section of Wittgenstein's *Remarks on the Foundations of Mathematics* IV, 48 and the linguistic analogue proffered by Wright:

But might it not be said that the *rules* lead this way, even if no-one went it? For that is what one would like to say - and here we see the mathematical machine, which, driven by the rules themselves, obeys only mathematical laws and not physical ones.

And analogously,

But might it not be said that this sentence is grammatical, and has the meaning it does, even if no-one considers it? For that is what one would like to say - and here we see the language-machine which, driven by the rules of the language themselves, obeys only linguistic laws and not physical ones (Wright, 1989: 238).

the work of Emil Post and the mathematisation of syntactic structure. I will briefly touch on this aspect of early generative grammar in the next section. The metatheory alluded to by Peregrin is captured in the field of formal language theory (FLT) which has since Chomsky's early work been dominated by generative (proof-theoretic) mathematics.

We do not have the space to enter into a protracted discussion of FLT here. A few details should suffice. FLT involves the mathematical characterisation of classes of formal languages. A formal language, in this sense, is a set of sequences of strings over a finite vocabulary. The members of this set vary according to the field to which we apply formal language theory, i.e. words if we are talking about natural languages or states if we are talking about programming languages etc. Furthermore, in formal language theory we are concerned with the finite ways in which these languages can be described, "FLT deals with formal languages (= sets of strings) that can be defined by finite means, even if the language itself is infinite" (Jäger and Rogers, 2012: 1957). This is usually done by means of formal grammars (i.e. sets of rules by which we construct well-formed sentences or answer the membership problem for a class of structures).⁴ The linguistic side of FLT is directed towards the goal of describing the various constructions of natural language syntax in terms of the class of formal languages in the infinite hierarchy of such languages that best captures it.⁵

Lastly, the concept of mathematisation is by no means idiosyncratic to the development of linguistic theory. In the natural sciences too, problems are often approached with a certain sort of mathematical transmogrification in practice (see Part II.5. for a discussion of targetless modelling). As early as Galileo, the natural world has been considered amenable to precise mathematical characterisation. Considering the observable phenomenon of free falling objects, a pattern emerges between distance and time (i.e. distance is proportional to time squared). It does not, however,

⁴In general rules look like $\alpha \rightarrow \beta$ where the arrow denotes "replace α with β " and α and β are sets of strings from the alphabet (either terminal or nonterminal).

G will be said to *generate* a string w consisting of symbols from Σ if and only if it is possible to start with S and produce w through some finite sequence of rule applications. The sequence of modified strings that proceeds from S to w is called a *derivation* of w . The set of all strings that G can generate is called the *language* of G , and is notated $\mathcal{L}(G)$ (Jäger and Rogers, 2012: 1957).

Another important component of formal language theory is decidability. Given a string w and a formal language $\mathcal{L}(G)$, there is a finite procedure for deciding whether $w \in \mathcal{L}(G)$, i.e. a Turing machine which outputs "yes" or "no" in finite time. In other words, a language $\mathcal{L}(G)$ is decidable if G is a decidable grammar. This is called the membership problem.

⁵"Generating" formal languages through grammars is not the only means of characterising or demarcating them. These languages have been shown to be susceptible to model-theoretic treatment in terms of monadic second order logic. For instance, Büchi's (1960) result showed that a set of strings forms a regular language if and only if it can be defined in the weak monadic second-order theory of the natural numbers with a successor. Thatcher and Wright (1968) showed that context-free languages "were all and only the sets of strings forming the yield of sets of finite trees definable in the weak monadic second-order theory of multiple successors" (Rogers, 1998b: 1117). I thank Geoff Pullum for clarifying some aspects of this alternative to me.

suffice to merely describe this pattern in plain language.

The mathematisation of the problem consists in our being able to specify the relation between distance and time in a precise way, a specification that is not possible using qualitative language. But note here that the relation between the *qualitative* concepts of distance and time plays an important role in what we call the ‘mathematisation’ of the problem [...] What is interesting, however, is that from the uses of mathematics as a type of tool for reconstruction emerges a representational framework with a life of its own (Morrison, 2015: 2-3).

Other examples include the technique of the thermodynamic limit in particle physics. In order to explain the breakdown of electromagnetic gauge invariance, physicists help themselves to the notion of a phase transition. Phase transitions involve a thermodynamic limit or “in other words, we need to assume that a system contains infinite particles in order to explain, understand, and make predictions about the behaviour of a real, finite system” (Morrison, 2015: 27). Similar techniques are used in population genetics in which the mathematisation of finite real populations results in models of infinite populations and their properties. In some of these cases, such as the field of mathematical physics, the line between mathematics and the natural world is irrevocably blurred.

One key difference, however, between mathematisation in the natural sciences and linguistics is that the mathematics is treated as a modelling tool in the latter and an explanandum in the former, especially in the case of infinity. We will delve into these issues in much greater detail throughout the thesis, with special attention paid to the posit of linguistic infinity. For now, let this serve as an introduction to the concept which will be developed through the grades of involvement in section 3. In the next section, we will move on to the mathematical approach which underlies Chomsky’s seminal *Syntactic Structures* and the methodological traces it left behind in the ensuing field of linguistics.

1.3 The Legacy of *Syntactic Structures*

Though often unappreciated, the work of Emil Post has had a profound effect on the field of linguistics.⁶ Chomsky’s *Syntactic Structures* (1957) (and other papers at the time such as his 1956b) offers a distinctly proof theoretic approach to the idea of a grammar or rule-based production system. A Post canonical production system is a just a tuple $\langle A, I, R \rangle$ with a finite vocabulary A or “axioms”, a set of initial words I (disjoint from A) and a finite set of transformations or production rules R (these are binary relations) (each of which has an antecedent x and consequent y such that $(x, y) \in R$ which ensures that there are no free variables in the consequent that are

⁶See Pullum 2011 for discussion.

not in the antecedent). This system resembles familiar natural deduction systems in propositional and predicate logic. If you want to prove a specific conjecture, you start with the members of A and derive the conjecture *via* repeated application of R . “In particular, [Post] developed a generative characterization of the recursively enumerable (r.e.) sets, and later laid the foundations of recursive function theory” (Pullum, 2011: 280). In so doing, he provided a formalism for modelling the concept of a logical proof. In a sense, this provides a proof procedure for discovering the strings or formal language “generated” by a given system of rules (on a finite alphabet or vocabulary). In addition, Post canonical systems are Turing complete i.e. belong to the same class as Turing machines.

It is not surprising that the mechanism of a Post production system became central to the concept of a generative grammar used in linguistics. Chomsky (1957: 22) defines $[\Sigma, F]$ grammars in the following way:

$[\Sigma, F]$ Grammar : Each such grammar is defined by a finite set Σ of initial strings and a finite set F of “instruction formulas” of the form $X \rightarrow Y$ interpreted: “rewrite X as Y ”.

From these (and earlier) insights, formal language theory was born. As we saw above, FLT is the abstract theory of syntax in which languages are viewed as sets of strings without semantic content. For Chomsky, a $[\Sigma, F]$ grammar (or generative grammar) is a system of rewrite rules on sets of terminal (or words) and non-terminal strings (or phrasal categories).⁷ *Syntactic Structures* suggests a proof that the syntax of natural language cannot be captured by a specific kind of formal language (a regular or finite-state language), although this result is technically never proven (and later developments such as Shieber (1985) prove that the more complex phrase-structure grammars are also inadequate for some languages). Natural language syntax is now assumed to be located somewhere in between the context-free and context-sensitive languages of the traditional ‘Chomsky Hierarchy’, specifically within the mildly-context sensitive class.⁸

A central insight, namely that the linguistic capacity of language users is unbounded, is what led Chomsky to develop the mathematical analogue of a computational system in order to represent this phenomenon. As Lobina (2014) notes, Chomsky adopted this position at a time when the terms “computation” and “recursion” were used interchangeably and this might explain his insistence of the centrality of recursion within linguistics itself. The idea that natural language syntax could

⁷There are some differences between Post canonical systems and phrase-structure grammars such as divergent notions of variables.

⁸A set of languages \mathcal{L} is mildly context-sensitive iff (1) \mathcal{L} contains all the context-free languages, (2) \mathcal{L} can describe the copy language and certain cross-serial dependencies of that sort (like in Swiss German, cf Shieber (1985)), (3) \mathcal{L} is parsable in polynomial time and (4) \mathcal{L} has the constant growth property.

be represented by recursive rule systems or as computational devices was not completely novel as we saw with Bar-Hillel.⁹ Chomsky's work however emphasised the need for greater precision which ultimately lead to linguists or mathematical linguists taking more interest in formal language theory, i.e. the mathematisation of syntax. However, mathematical linguists tend to restrict themselves to linguistically interesting or motivated investigation into the infinite classes of formal languages and their respective complexities. For instance, context-sensitive languages and regular languages are generally of little interest given that it is unlikely that natural language constructions can be found within these formal parameters (oversimplifying, the latter are too restrictive and the former too complex).

Whereas in formal language theory the use of sets can be viewed as an abstraction or convenience in some sense, within recursion theory sets (in terms of functions) are not optional. If recursion is assumed to be a feature of the natural landscape as opposed to merely a feature of our models, we move toward a more pervasive mathematisation.¹⁰

It is not my purpose to detail the developments of formal language theory or formal syntax here. The claim, which I hope to impress upon the reader, is that the beginnings of the generative or biolinguistic movement have nontrivial logical and mathematical foundations. The historiography of linguistics is a much more complex matter than I have shown in the previous two sections, where my intention was illustrative rather than comprehensive. Needless to say, Carnap, Goodman, Quine, Harris all deserve mention within a more complete story. However, such a task is beyond the scope of the present work (see Newmeyer (1996) for a generative approach and Tomalin (2007) for a more objective attempt).

It is important to mention, at this juncture, that the generative tradition in linguistics offered an approach to central aspects of natural language comprehension and production that were previously unaccounted for, namely productivity, learnability and creativity. The tools of this restricted class of Post canonical systems, i.e. generative grammars, allegedly provided insight into something that could not be approached without this mathematical apparatus, specifically the question of how a finite system could generate an infinite output (a key aspect of productivity, learnability and creativity on this account). Chomsky (2000) claims that the explanation of

⁹The history of the term "recursion" in linguistics is extremely messy. It is not my intention to be embroiled in that controversy here. The core idea is that recursive functions introduce a property of self-reference. This usually involves two steps. One which specifies the condition of termination of the recursion or the base case and the recursive step which reduces all other cases to the base. The Fibonacci sequence is an example of this procedure, so too is Bar-Hillel's (1953) definition of a French sentence.

¹⁰Notice, even Chomsky's famous (1956b) disavowal of the relevance of stochastic grammar formalisms, in which approximation through continuous mathematics is the goal, can be seen as motivated by mathematisation. The statistical methods of continuous mathematics do not generally make a mathematical object of the target domain but rather treat it as physical process capable of "approximate" characterisation.

natural language creativity only became available with the advent of computability theory in the 20th century (and before then seemed like a contradictory property for a physical system to possess), or “[a]dvances in the formal sciences provided that understanding, making it feasible to deal with the problems constructively” (Chomsky, 1995: 4). Again, recursion commanded a central explanatory role in linguistic theory. However, Lobina (2014) claims that with the emergence of “merge” this role somewhat diminished.¹¹

To be sure, the early years of generative grammar saw the employment of Post’s production systems, a more obvious recursive formalism, but this aspect of the theory hasn’t translated much with the advent of merge, given that a recent description delineates it in very general terms as a set-theoretic operation in which repeated applications over one element yield a potentially infinite set of structures, drawing an analogy between the way merge applies and the successor function (Lobina, 2014: 2).

Despite the shift in mathematical features of the field, the analogy with the natural numbers persists throughout the discipline. Consider this claim made by Pinker (1994: 86) in a popular book on linguistics.

By the same logic that shows that there are an infinite number of integers – if you ever think you have the largest integer, just add 1 to it and you will have another – there must be an infinite number of sentences.

Similarly, Hinzen and Uriagereka (2006) draw even stronger conclusions concerning the connection between linguistics and mathematics, well-within the biolinguistic tradition.

[T]he human language faculty poses much the same explanatory problems for contemporary physicalism as the mathematical faculty does (72).

and,

Chomsky’s technical correlation between language and mathematics is also well-taken, given the biological isolation they both share as systems instantiating discrete infinity. However, if the latter can be abstracted from FL [faculty of language], just how much of mathematics is FL using to begin with? [...] Surely the successor function fits naturally into syntagmatics (84).

We shall return to this issue below, but for now suffice to say that the analogy between linguistics and mathematics goes deeper than just that of the merge operation and the successor function. The ensuing movement only added to the mathematical

¹¹Again, these issues are too unclear to delve into here. In some sense Merge can be considered “recursive” in that it is a general procedure that can apply to its own output, in another sense it can be understood as producing constituents of types that contain other constituents of the same type.

foundations of linguistics. The use of sets and functions have become ubiquitous in linguistic theory but the mathematical apparatus does not stop at these relatively benign tools of characterisation or formalisation. Along the way, merge and recursion somehow became in need of physical interpretation.

On the other side of the fence, semantics was also undergoing mathematisation. Following the pioneering work of Lambek (1958), Montague famously stated that he “reject(s) the contention that an important theoretical difference exists between formal and natural language” (1970a: 188). In fact, this is a point (reiterated below) upon which he believed himself to be in agreement with Chomsky as evinced in ‘Universal Grammar’.

There is in my opinion no important theoretical difference between natural languages and the artificial languages of logicians; indeed, I consider it possible to comprehend the syntax and semantics of both kinds of language within a single natural and mathematically precise theory. On this point I differ from a number of philosophers, but agree, I believe, with Chomsky and his associates. (Montague, 1970b [1976]: 222).

Indeed, according to Thomason (1974), Montague held syntax and semantics to be branches of mathematics. It is not clear that Chomsky and his associates explicitly held the view Montague attributes to them. Nevertheless, the idea which does seem attributable to generative linguists is that structures of natural language not only lend themselves to mathematical characterisation but moreover there is some “special” connection between the faculty of language and that of mathematics as they both “instantiate” a mathematical property, namely discrete infinity.

All approaches agree that a core property of FLN is recursion, attributed to narrow syntax in the conception [...] This capacity of FLN yields discrete infinity (a property that also characterizes the natural numbers) (Chomsky, Hauser and Fitch, 2002: 1571).

This aforementioned connection gets to the heart of the distinction between *formalisation* and *mathematisation* which I propose and hope to illuminate in the subsequent sections.

Whereas *Syntactic Structures* could be read as an application of discrete mathematics to language, opening up new vistas in the study of formal grammar, *Aspects* made clear that formalisation, and indeed the study of languages/grammars was not an end in itself. The goal was rather to reverse-engineer the structure of language to discover the mind that made it possible in the first place (Boeckx, 2015: 128).

Thus, if all of the grammars of natural language imply a discrete infinity of expressions, then this capacity is attributed to the language faculty by “reverse-engineering”. Statements like these are the hallmark of mathematisation. We will return to these issues in sections 2.1 and 3.2 respectively.

The next few sections delve into the details of the specific foundational projects in linguistics before moving to the grades of mathematical involvement compatible with these frameworks.

Chapter 2

Methodology and Ontology

As previously mentioned, I will argue that methodology and ontology can and do come apart in the foundations of linguistics. However, the ontological foundations of the science have traditionally had an influence on the “working linguist”. For instance, pre-Chomskyan structuralists are often characterised as nominalistically motivated. They made no assumptions about a universal grammar (UG) connecting the world’s languages or deep structure underlying surface forms and their categorisations and linguistic descriptions were often restricted to the corpora of specific languages, i.e. natural language tokens.¹ Although, certain proponents of this framework, such as Hockett, viewed linguistics as an essentially “classificatory science”, this methodology by itself did not entail that natural languages were not abstract objects or mental ones. Nevertheless, the view has been thoroughly supplanted by the generative tradition and will thus not be the focus of the current study (even though some of its core insights have been retained within computational linguistics, see Lenci 2008). In addition, it incorporated a physicalist view which is included in conceptualism (as a proper subset) and the nominalism at its core will be discussed in Part III. In what follows, I will briefly outline some of the key aspects of the generative movement in terms of its methodology and ontology and then attempt a similar characterisation of the linguistic Platonist programme before considering a pluralist position which aims to unite both views *inter alia*.²

¹“Non-linguists (unless they happen to be physicalists) constantly forget that a speaker is making noise, and credit him, instead, with the possession of impalpable ‘ideas’. It remains for linguists to show, in detail, that the speaker has no ‘ideas’, and that the noise is sufficient” (Bloomfield, 1936: 23).

²The various views connected under the banner of “public language” also deserve mention as an ontological basis of natural language. These views tend to agree on the social normative nature of natural language conventions (although, they disagree on exactly what this nature looks like). However, the public language theorists tend not to have a clear (or unified) linguistic methodological approach, especially when it comes to syntax. Thus the methodology-ontology question is mute in this case. For an exception, see Peregrin 2008 for a discussion of the compatibility of inferentialism and formal semantics.

2.1 Linguistic Conceptualism

Linguistics as Psychology

Before any serious overview can be presented on this subject, some terminology has to be settled. I have here opted for the term “conceptualism” partly to avoid thorny (often times ideological) issues of the alternatives such as “biolinguistics” and “generative grammar”.³ The exact nature of the biolinguistic programme in linguistics is not always clearly defined (as with many other disciplines). The questions of its relationship to previous incarnations of the generative tradition and the question of the relationship between biolinguistics and generative grammar itself are surely important ones. These questions will not receive thorough treatment here (see Lasnik 2005, Boeckx 2015 for discussion). I myself consider the biolinguistic position to be stronger in its claims and link to biology than the position represented by generative grammar which is usually understood as concerned with language and its relation to the cognitive system. For instance, issues related to language evolution are more central to the biolinguistic approach than to generative grammar, hence the recent discussion of merge as an evolutionary mutation of some sort. Another way of putting the point is in terms of the goals of linguistic theory. In *Aspects of the theory of Syntax* (1965), Chomsky differentiates between three nested kinds of adequacy conditions for a theory of grammar, each more inclusive than the last. The three related linguistic desiderata are (1) observable linguistic performance, (2) native speaker judgements and (3) language acquisition. The first is the class of *observationally adequate grammars* which are those grammars which only account for corpora or observed utterances of speech. Naturally, these do not give us much traction on (2) and (3). Chomsky then suggests a class of *descriptively adequate grammars* (DAG) which aim to capture the psychological facts of native speaker intuitions, thereby addressing (1) and (2). However, these latter grammars are inadequate on count (3) and thus require us to ascend to the level of *explanatorily adequate grammars* (such as can be found in the Principles and Parameters framework).

In later developments, linguistics moves further afield from these levels. Within the minimalist programme, Chomsky identifies another level, one ‘beyond explanatory adequacy’ which he calls ‘natural adequacy’. Thus, the goal of linguistic theory moves away from the concerns which characterised generative grammar of merely describing linguistic intuitions or language acquisition to a theory of language as a natural object. A natural object is explained as something being bound by the biological and physical universe, as opposed to the mathematical and conventional ones.

In principle, then, we can seek a level of explanation deeper than explanatory

³In Part II, I will offer an alternative means of capturing the scientific import of the latter term. I could have chosen “mentalism” just as well.

adequacy, asking not only what the properties of language are, but why they are that way (Chomsky, 2004a).

What I am interested in is the philosophical position espoused by the view as to the nature of natural language and the corresponding focus of linguistic theory, i.e. the common conceptualism (or mentalism) at the base of both biolinguistics and generative grammar. In this way, I hope to understand the more formal aspects of the paradigm and how its claims concerning recursion, discrete infinity and merge fit into this naturalistic picture.

On the conceptualist view, linguistics is a branch of cognitive psychology and the foundations of *linguistics proper* (a term introduced in Katz and Postal 1991 to describe the job of the working linguist) are eventually to be subsumed by biology or neuroscience. A more general understanding is that linguistics constitutes the study of a biological system which is responsible for language generation. Therefore, it incorporates a physicalist ontology. However, there are several important differences between this picture and the American structuralist movement (which proffered a similar albeit more restricted ontology).

The first important difference between this programme and its predecessor is that the linguistic universe is expanded to include more than just the observable physical utterances of speakers (desideratum (1) above) but also the innate linguistic structures in their minds/brains. Early generative grammar postulated “deep structures” which through various transformation rules produced the surface forms of sentences which we can perceive directly. These deep structures were just like the expressions generated by the recursive rules of a context-free grammar or a phrase structure grammar (i.e. a restricted version of a Post canonical system), such as $S \rightarrow NP, VP$; $VP \rightarrow Aux, V, NP$ etc. Surface forms are then linked to this structure *via* various “transformations”. Although Chomsky would eventually eschew the notion of ‘deep structure’ in ‘the Minimalist Program’ (1995), the idea that linguistics is concerned with hidden cognitive structures was important progress for the field and a major departure from its predecessors. Katz (1971) refers to it as a “second linguistic turn” akin to the Democritean revolution in early scientific thought.

I-Language and Grammar

So what kind of thing is a language on this view? For Chomsky it is psychological in nature, linked to the state of the so-called “language faculty”.

We can take a language to be nothing other than a state of the language faculty [...] So let’s take a language to be (say, Hindi or English or Swahili) a particular state attained by the language faculty. And to say that somebody knows a language, or has a language, is simply to say their language faculty is in that state (Chomsky, 2000b: 8).

Furthermore, although different languages can be described in terms of different states, there are general, “universal” and innate properties underlying all of these seemingly distinct languages, so much so that “deep down, there is only one human language” (Chomsky 1995: 131). This is the infamous Universal Grammar or UG hypothesis or in other words “there are aspects of the linguistic system acquired by the child that do not depend on input data [...] Some cases of this type, it has been argued, reflect the influence of a genetically prespecified body of knowledge about human language” (Pesetsky, 2009). Interesting though this may be, the status of the UG hypothesis will not be my concern in this thesis.

More importantly, a grammar is supposed to be a scientific theory of the mental behaviour or of the state that the language faculty is in (so different states for English, Mandarin, isiZulu etc.).

A better usage would be to restrict the term “grammar” to the theory of language, and to understand the language as what we may call “I-language” where “I” is to suggest “intensional” and “internalized” (Chomsky, 1990: 678).

The grammar thus describes a language which is “internalised” in the sense of being located in the mind and eventually the brain of the ideal speaker-listener and “intensional” in terms of a function that determines a restricted set of expressions or a “grammar” and not the entire discretely infinite faculty of language itself. Another word for the former characteristic is “individualistic” or non-relational in the sense of possessing properties that depend on or are related to only internal mental features of the language-user.

The data associated with grammar construction is the (oft-criticised) introspective judgements of native speakers or the linguists themselves. This is the idea that our ‘intuitive’ linguistic judgements can serve as the “voice of competence” and thus the primary data for linguistic theories. Marantz (2005) argues that the role of introspective data has been mischaracterised and linguists’ judgements stand as “proxies” or meta-data aimed at representing and not reporting the data. So the judgements of linguists’ merely indicates the need for further corpora based or distributional investigation for the sake of confirmation. In addition, psycholinguistic and neuroscientific techniques are being used by linguists not only for the sake of testing theoretical hypotheses (as was the case with initial psycholinguistic research) but also for the sake of forming new hypotheses. Using the various techniques of psychology and neuroscience establishes greater compatibility with the other cognitive sciences and would presumably be more in keeping with the naturalistic or scientific turn which the field has undergone under the conceptualist framework. Despite these advantages and larger integration within cognitive science, many linguists have resisted the incorporation of scientific techniques and insisted on the formal characterisations in terms of proof theory initially and now more set-theoretic techniques.⁴

⁴See Jackendoff 2002 and 2007 for an attempt and reflection respectively of the incorporation of generative linguistics within the larger cognitive scientific movement.

Formal Aspects of the Theory of Syntax

The question which most concerns this research is that of the formal aspects of the theory, such as recursion in early generative grammar and merge in minimalism, and their relation to the mathematical realm. Whether through recursion or merge, conceptualism seems to advocate for an extra-biological, or at least special, claim that language is somehow mathematically unique. This claim starts with the idea that the human language capacity (and cognitive capacity in general) is supposed to be understood as finite in its resources, yet one apparent aspect of natural language is its creative nature, assumed to be capable of (discretely) infinite expression.

The most striking aspect of linguistic competence is what we may call the ‘creativity of language’, that is, the speaker’s ability to produce new sentences that are immediately understood by other speakers although they bear no physical resemblance to sentences that are ‘familiar’ (Chomsky, 1966: 74).

Most linguistics textbooks start with the claim that natural language is infinite (for examples, see Lasnik 2000, Sag *et al* 2003, Yang 2006). Some linguists even go as far as to claim that infinity is the only linguistic universal (Epstein and Hornstein 2005). This aspect of the tradition has led to some criticism (see Pullum and Scholz 2010) but it has also led to the connections with the realm of mathematics as we saw in section 1.3 and can glean from these rather speculative comments in Chomsky (2010: 48):

The “gigantic development of the mathematical capacity is wholly unexplained by the theory of natural selection, and must be due to some altogether distinct cause,” if only because it remained unused. One possibility is that it is derivative from language. It is not hard to show that if the lexicon is reduced to a single element, then unbounded Merge will yield arithmetic.

In ‘On Phases’, Chomsky is more explicit on how this procedure is to be accomplished. Even though as Tomalin (2007: 1795) notes, if the lexicon contains a single element, then merge cannot be applied without some sort of indexation.

Suppose that a language has the simplest possible lexicon: just one LI [lexical item], call it “one”. Application of Merge to that LI yields $\{one\}$, call it “two”. Application of Merge to $\{one\}$ yields $\{\{one\}\}$, call it “three”. Etc. In effect, Merge applied in this manner yields the successor function. It is straightforward to define addition in terms of Merge(X, Y), and in familiar ways, the rest of arithmetic (2005: 6).

Talk of recursion, arithmetic, discrete infinity and the set-theoretic operation of merge seems to suggest a deeper analogy with the formal sciences as mentioned in the previous section. In a lecture in 2011 at Carleton University, Chomsky claimed that “perhaps the most elementary property of human language is that it consists of

a discrete infinity of interpretable expressions – so there's five-word sentences, and six-word sentences, no five-and-a-half words sentence, so it goes on indefinitely like the integers. That's kind of unusual, there's nothing like that known in the biological world." In my view, it is this alleged divergence from other aspects of the natural or biological world which leads the linguist from a purely naturalistic endeavour to a partly formal one. The problem is that discrete infinity requires the apparatus of discrete mathematics for characterisation. This is not mere formalisation but rather a mathematisation of an object (or state) claimed to be biological in nature.

Let us consider the "merge" postulate for a moment. MERGE is an arbitrary operation on sets of syntactic items, essentially it takes two objects and combines them into one (labelled) object. It is *internal merge* which performs the role which recursion performed previously. *External merge* takes two distinct objects as input and *internal merge* allows embedding and thus allows for recursion. Furthermore, *internal merge* involves duplicating items within the operation. For instance, if we merge syntactic objects α and β to form the unordered set $\{\alpha, \beta\}$ and there is a γ such that $\gamma \in \alpha$ and we merge this object with $\{\alpha, \beta\}$, we would have two copies of γ in the resulting structure. In this way, we can account for all movement with minimal operations in the syntax (and various constraints on the operations). It is in explaining the "arbitrariness" of merge that we once again see a parallel with arithmetic.

Within the framework just outlined, there is also no meaningful question as to why one numeration is formed rather than another – or rather than none, so that we have silence. That would be like asking that a theory of some formal operation on integers – say, addition – explain why some integers are added together rather than others, or none (Chomsky, 1995: 208).

Here we are dealing with an arbitrary set-theoretic function which yields a discrete infinity of natural language expressions, i.e. biological output. In the case of arithmetic, the output of the successor function is not usually considered to be physical in any strict sense (although nominalists in the philosophy of mathematics might disagree). And yet both language and arithmetic are alleged to have sprung from the same well. Whether it is Turing machines, discrete infinity or merge, my claim is that the conceptualist approach involves a level of mathematisation of the "natural object" of language. Linguistics thus seems to view itself as a "special" science in a sense divorced from other empirical sciences, and wedded to aspects of the formal sciences. As I suggested in the first section, the move seemed to involve a departure from mere *formalisation* to a distinctive *mathematisation* of recursive elements specifically.

[I]n the earliest work, although recursive components were considered useful formal procedures that simplified the basic analytical framework, no strong claims were made concerning their biological status. Gradually, though, as the theory of GG developed [...] the role of recursion within the GG framework began

to acquire cognitive connotations, with the eventual result that [...] it has been hypothesised that recursion is a genetically-embedded computational procedure (Tomalin, 2007: 1785).

Given the conceptualist and more so the biolinguistic agenda (of integrating the study of language with other biological systems), the onus is on conceptualist to provide an evolutionary story for how such an admittedly “extra-biological” discrete infinity creating operation such as merge emerged from the physical world. In an attempt to do so, Hauser, Chomsky and Fitch (2002) go as far as to state that recursion or merge is the core property of the faculty of language (narrowly construed). They go further to state that merge is an evolutionary mutation which gave rise to human linguistic abilities (and perhaps arithmetic ones too). This story has met with large-scale criticism and an in-depth discussion of this postulate is beyond the current scope but it does serve as an example of the centrality of the recursion and discrete infinity claim within the biolinguistic paradigm. At the very least, biolinguists seem to have to account for the role of such set-theoretic operations in their science which is alleged to be an empirical one. Is talk of recursion and infinity merely descriptive or modelling of a physical feature of our biological makeup or is it to be taken more literally as an actual feature of linguistic knowledge and generation? And if so, how? In section 3.1. I will offer such an explanation on their behalf but for various reasons they may be less inclined to adopt it. In 3.2 I will discuss these issues in somewhat more detail.

2.2 Linguistic Platonism

We saw a framework in two different spirits in the previous section. On the one hand, the conceptualist movement aims to offer a naturalistic account of the study of natural language, one which draws on aspects of human psychology and (eventually) biology. On the other hand, it places the explanation of the discretely infinite capabilities of natural language users at the forefront of its scientific agenda thereby relying heavily (almost exclusively) on set-theoretic and other formal methods of exposition.

Abstract Ontology

Platonism begins with a critique of this apparent bi-polar methodological aim and its resulting “incoherent” ontology as proponents of this view claim.

Platonism is an account of linguistic foundations which holds that linguistics is the study of abstract mind-independent objects. The Platonist takes all of the syntactic and semantic structure posited by grammars not merely as useful tools for describing mental states but actual features of an objective linguistic reality. A natural language, like a formal language, is an abstract object. In a sense this is an externalistic conception (i.e. E-language, or P-language as it is called in Chomsky 1990) since

natural language is claimed not be defined within human biology (or psychology), but importantly unlike E-languages this view holds that natural language also exceeds the bounds of physicalism. Languages on this view are of the same ontological type as sets and numbers, thus not located in space or time (or springs from the same *a priori* “faculty of intuition” according to Katz (1981)). Therefore, linguistics has to be a formal science on par with mathematics. If conceptualism involved a partial mathematisation of the object of inquiry, then Platonism is the case of full mathematisation. There are three main components of this view (which I think are present in its various statements and restatements), (1) an epistemic claim, (2) an ontological claim and, (3) a methodological one. I will (very) briefly describe each in turn (I will hold off on a critique until section 3.3).

We will start with (1). The claim here is that the conceptualist movement confuses (i.e. identifies) the knowledge we have of a language with the language we have knowledge of, *contra* the nature of the dyadic “knowledge relation”. Simplistically, if linguistics is the study of linguistic knowledge or linguistic competence (i.e. an I-language), then this should presuppose the existence of the language independently of the knowledge of it.

[F]or any domain X, knowledge of X systematically depends on the existence of X. For some real rabbit, R, one can determine in principle R’s average blood pressure. There can be knowledge of that average blood pressure. But a question about the average blood pressure of the Easter Bunny has no answer (Postal, 2003: 235).

If knowledge is a two-place predicate or relation (as is assumed on this view) then the *relata* have to be existent to make sense of any knowledge claim. Devitt and Sterelny (1989) take up issue with this epistemic identification likewise. Although, they question in what sense a language user can be said to “know” her language. They proffer a few candidates for what this knowledge relation could constitute, each more problematic than the last (we just do not seem to have the right kind of epistemic access to the rules of our language, the rules which the grammar posits). Eventually they adopt a know-how (as opposed to a propositional or know-that) account of linguistic knowledge. A similar approach is suggested by proponents of dynamic syntax (Cann *et al*, 2012).

Claim (2) asserts that there is a deep ontological incoherence at the root of the generative approach. The claims of the theory are empirical and mentalistic, yet the approach is formal and mathematical. There are quite a few arguments offered in favour of a consistent Platonistic ontology in its stead. I will mention two such arguments here, connected to the issue of mathematics. One of which might be more familiar than the other.

The received view claims that an NL is something psychological/biological, in the baldest terms, a state of an organ, that aspect of the brain that permits

NL to arise [...] And yet it has been unvaryingly claimed in the same tradition at issue that NL is somehow infinite. These two views are not consistent (Postal, 2003: 242).

The first major argument in favour of Platonism is based on the ontology of sentences. Following a Quinian type/token distinct, this argument takes the form of a *modus tollens* in Postal (2009). If NLs are psychological or biological (I-languages, states of the mind/brain etc.) then sentences of NLs are also physical. But sentences are not at the correct level of abstraction to be aspects of our brains, rather if there is anything to this conceptualist claim, we have to be talking about sentence tokens. Sentence tokens, utterances or expressions, whatever the preferred term, are indeed physical objects but they are clearly not the objects of linguistic study (perhaps they are appropriate units of study for psycholinguistics?). However, sentences, unlike their tokens, are not to be found in the physical world. In other words, we can ask when and where a sentence was tokened (uttered or expressed) but not when and where the Spanish sentence *¿Quién engañó a Roger Rabbit?* is or was.

Another argument in favour of a type/token distinction here could come from the philosophy of science. Grammars, taken as scientific theories, can only be understood as structural descriptions of sentences *qua* types not tokens for any theoretical generalisations to make sense. Otherwise we would be left in the same state as the American structuralists who assumed no generalisations between languages or even the specific corpora under description (or so the caricature goes). I think this line might be a bit misguided though. For instance, a biologist might “abstract” over specific tokens of animals and speak of “species” or types of animals without requiring those types to be non-spatio-temporally extended. Is biology a formal science? Probably not.

Nevertheless, I do think that the problem above can be blocked or rather it is not a viable response to Postal’s case against biolinguistics. This is the case because the response is incompatible with the core feature of natural language posited by the Chomskyan paradigm, namely discrete infinity. If we were to collect all of the sentence tokens of every language throughout their individual histories (we can even throw in the extinct languages), we would not arrive at a denumerable infinity, i.e. a mapping from the sentence tokens to the natural numbers. The assumption is that there is no analogue in the biological world (see Chomsky quote at the end of the previous section and Hinzen and Uriagereka (2006)). Put in another way, type-talk might be a useful convenience for theorising in biology (and other sciences) but it is a necessity for linguistics since linguistics is precisely about sentences types themselves notwithstanding the issue of their instantiation in the real world. If sentences were sets (or set-like) and linguistics was a formal science, these worries would presumably disappear.⁵

⁵This is not necessarily the case. There is still no analogue of the Peano axioms or an inductive

How Big is NL?

There is a more technical (and related) argument in favour of Platonism which argues for the nondenumerable infinity of natural languages (i.e. strictly greater than the set of natural numbers). This possibility attacks the heart of the competence model of grammar. Although it has been largely neglected in the literature, I think it is pertinent to the issue at hand since it involves explicit mathematisation. According to this result, the cardinality of natural language is not even a set but a “megacollection” (or proper class). If this is the case, then the Post canonical system procedure or the generative grammar one is mistaken and does not determine or enumerate the set of natural language sentences since sets have fixed cardinalities. Thus, these procedures cannot capture the magnitude of natural language.

It is not possible to reproduce this proof in any detail here (I refer the reader to Langendoen and Postal (1984) or for a shorter overview in Katz (1985)). I will, however, outline the strategy of the proof.

Briefly, we start by defining the concept of coordinate compounding. That is, any constituent is a coordinate compound of a given grammatical category iff (1) it is part of that grammatical category, (2) it has two immediate constituents itself and finally (3) these constituents are conjuncts.

From this they draw the claim that all sets of constituents are closed under coordinate compounding (and thus contain their own coordinate projections⁶) and the same applies to constituents of all sentences. Consider the sentence *I exist* which can be extended to another sentence *I know that I exist* and *I know that I know that I exist* etc. Call the set of all of these sentences X . We know that (and Langendoen and Postal admit that) X is countably infinite. If X is closed under coordinate compounding then a continuum is created since X contains X' which contains all the elements of X and their coordinate projections, i.e. *I exist and I know that I exist and I know that I know that I exist...*

The reason is that each coordinate compound sentence of X' can be put in one-to-one correspondence with a member of the power set of X .⁷ Therefore, since the size of a power set of a countably infinite set is \aleph_1 or 2^{\aleph_0} , X' is also of this cardinality. But since $X' \in X$, X must at least be \aleph_1 . Importantly, “at no point can a set of sentences be obtained that exhausts an NL having sentence coordination governed by the closure law” (Langendoen and Postal, 1985: 58). This leads us straight into the Vastness proof. The proof proceeds by contradiction. You start by assuming that some NL is a set Y (i.e. has a fixed cardinality) and then show that Y and its

proof for the discrete infinity of natural languages. When infinity is claimed, it is assumed. See Pullum and Scholz (2010) for details.

⁶These are like functions that take constituents of the same category and outputs a set of conjoined items, e.g. take a set of items $\{Bob, Susie, Felix\}$, then the coordinate projection of this set is *Bob, Susie and Felix*.

⁷Non-compound $s \in X' \mapsto \{s\} \in X$, coordinate compound $s \in X' \mapsto$ projection set...

coordinate compounds are contained in the original set and thus of less cardinality than Y but you also show (*via* Cantor's theorem) that the projection sets (of the compounds) are equinumerous to the power-sets of Y , thus of greater cardinality than Y . Contradiction. Therefore, NL is not a set.

The immediate and devastating consequence of this theorem is supposed to be that all generative (constructive or proof-theoretic) grammars are rendered useless in characterising NL from the onset since they assume a denumerable infinity. As a result, a large research paradigm in linguistics should be abandoned. Non-constructive grammars would have to supplant the generative methods.⁸ In addition, grammars cannot be about physical brain-states as *per* the biolinguistic paradigm since this would be mathematically impossible. Brain-states are finite (whatever this means exactly) and without generative procedures for arriving at the requisite infinity (since these max out at \aleph_0), unlikely objects of our grammars, i.e. generative grammars can never describe linguistic reality fully. Put in another way, generative grammars supposedly provide us with a bridge from the finite to the infinite in linguistics. If the full complexity of natural languages exceeds the reach of generative procedures, then linguistic competence cannot be identified with the target of grammatical theory (conceived of as concerning natural language) since competence is only a proper subset of natural language (this would also help with the epistemic claim).

Mathematical Methods

Finally, moving onto (3) or the methodology of linguistic Platonism, this section has attested to the possibility of the mathematical nature of the study of natural language. We are now in a position to appreciate the depths of the methodology of the Platonist programme. Although, for instance, Katz (1985) insists that the working linguist would be unaffected by this shift in foundational interpretation and rather it is more compatible with her quotidian task, I think that Platonism marks a further departure in terms of methodology. On the extreme side, the Vastness proof would allow for elements of the structures of linguistic reality to be directly proven *a priori* (a possibility Fodor 1981 refers to as "The Wrong View"). More importantly, our linguistic intuitions should serve us well in grammar construction and thus the need for proxies (à la Marantz above) or corpora studies is rendered otiose, since mathematicians do not need to pool mathematical intuitions for data. In addition, the contemporary trend towards the inclusion of linguistics within the larger cognitive scientific framework would also be blocked by this foundational picture, since the methods of the latter would not be necessary to shed light on the former.

On this view, mathematical and linguistic reality would be amenable to the same methodological treatment. Ontologically speaking, linguistic objects (sentences, nat-

⁸See Pullum and Scholz (2001) for a discussion of the cardinality neutrality of model-theoretic approaches to grammar. This feature will become important in section 4.1.

ural languages) are like mathematical objects, such as sets, groups and numbers, and thus beyond the physical realm.

2.3 Linguistic Pluralism

There is a neglected position on the foundations of linguistics which is directly informed by the alleged heterogeneity of linguistic objects and methodology. This view embraces the so-called “hybrid” nature of generative grammar and offers a thoroughgoing pluralism as its genesis in both ontology and methodology. It embraces both the conceptualism and Platonism of the sections above (among other positions).

From Plural Methods to Hybrid Objects

The chief proponent of this view is Robert Stainton (1996, 2001, 2006, 2014). I will provide an overview here and then discuss briefly why I do not believe that this alternative gets to the heart of the issues which form the target of the present work. I am breaking the expository tradition I established in the previous two sections here. The reason is that arguing against pluralism allows us a natural point to justify many of the assumptions of the discussion so far and the presuppositions of the thesis more generally. This early justificatory project should serve us well in later sections and parts for either the clarification of confusion or the preemption of objection.

Firstly, pluralism rejects the idea that there is any real distinction to be had with the concept of “linguistics proper”, which was coined to isolate syntax and semantics as the only properly linguistic domains. From the pluralist perspective, phonology, phonetics (and presumably pragmatics) are equally linguistic domains worthy of inclusion within any debate concerning the foundations of the subject. Naturally, once we move toward a broader methodological base, the ontological plurality seems to follow. In the spirit of this diversification of the properly linguistic, Stainton describes the corresponding metaphysical attitude in the following way.

My own view [...] is that natural language, the subject matter of linguistics, have, by equal measures, concrete, physical, mental, abstract, and social facets. The same holds for words and sentences. They are metaphysical hybrids (2014: 5).

Stainton provides various arguments based on some extremely interesting evidence to support this ontological attitude. I divine two main lines of argument for the pluralist claim. The first is that the ontological attitudes of the previous sections (with the inclusion of public language views and naive physicalism of the pre-Chomskians and Quine etc.) all have something to contribute to the subject matter of linguistics and therefore its corollary that each ontology misses out on something

essential about the picture individually. The second line is that on a specific reading of these ontological commitments, and contrary to *prima facie* impressions, these ontologies are perfectly compatible with one another.

Let us consider some of the evidence for the first claim, i.e. that physicalism, mentalism, Platonism etc. all have some important *sine qua non* piece of the linguistic puzzle to contribute. Above, I considered mentalism to be part of the physicalist persuasion. For the sake of distinguishing the ontological perspectives associated with Bloomfieldian physicalism and Chomskyan mentalism, I will follow Stainton in referring to the former as physicalism and the latter as conceptualism (as *per* my usage). One argument for the necessity of the physicalist contribution comes from phonetics, another from pragmatics. The argument from phonetics simply involves the truism that vocal and auditory organs, and the sounds which they produce, play a role in language production and comprehension. Phonetics concerns the physical movement of the vocal tract, the tongue, aspiration etc. These phenomena are clear candidates for physical aspects of natural language. On the other side, pragmatic phenomena seem to reference particular situations of interpretation. Interpreting indexicals such as *here* and *now*, deixis, demonstratives and other contextual elements point to the need to incorporate pragmatic aspects into the ontology of natural languages.

The cases for the inclusion of the conceptualist and Platonist ontologies have already been covered above. The point is that all of these facets contribute to the “properly linguistic” and the exclusion of any of them is tantamount to incomplete characterisation.⁹

Problems for the View

Despite the ecumenical spirit of the approach, there are a few aspects of this line of reasoning which I take to be questionable. Firstly, there is a distinction between the fact that a contribution is made by a set of phenomena to the description of a general (super)phenomenon and whether or not it counts as detrimental to abstract away from it. The idea of “linguistics proper” is an abstraction in my view. In the process of scientific investigation abstraction is a necessary tool. In many cases this process involves the omission of potentially connected or relevant material. The subject of Part II is concerned with idealisation and abstraction in generative linguistics. I will not preempt that discussion here except to say that the aim of such abstraction is often the isolation of the core or minimal structures responsible for a given phenomenon.¹⁰

⁹I abstract over the necessity claim of the social aspects of natural language which involves phatics and other similar speech acts. See Millikan (2005) and Peregrin (2015) for two distinct approaches to the legitimacy of public language as a object of scientific inquiry. See also Part II.6 for a discussion of the latter view within the context of formal semantics.

¹⁰Even Stainton’s frequent parlance of a certain phenomenon not “being exhausted” by a certain description or methodological characterisation is misleading as scientific investigation might not be interested in “exhausting” descriptions of phenomena but rather minimally representing them. The

Consider the case of phonetics. It is undeniable that most of the world's languages involve sounds in terms of combinations of vowels and consonants. Some of these combinations are quite complex, such as the click sounds of my native South Africa.¹¹ Nevertheless, they are not essential to an understanding of natural language *simpliciter* since there are languages which do not produce sounds at all. Sign languages operate on manual communication with physical gestures and signals to convey information. Some sign languages include haptic cues (used in communities whose members suffer from both deafness and blindness). They have syntax and semantics. They do not have phonetics or phonology. Assuming that phonetics contributes some essential aspect of natural language without which we do not have a characterisation of natural language relegates sign languages to the camp of the non-natural. Abstracting away from phonetics could therefore be considered benign.

A different rationale is required for abstracting away from pragmatic (and social) phenomena. One way of thinking about pragmatics is that it concerns language use "in context". But for anything to be used in context or in a context is to suggest that it had a prior form. A natural language can thus be thought of as a package comprised of syntax and semantics which is applied to real world communicative environments. For instance, a grammatical sentence can be attributed with a meaning based on the compositional rules of the language. The standard way of viewing the product of this procedure is that the sentence is endowed with a *literal meaning*. Now this sentence can be used to convey a myriad of other meanings depending on the circumstances (sarcasm, deception, play etc.). There might be pragmatic "rules" (such as Grice's maxims and some more complex tools of formal pragmatics) but these rules tend to be violable and imprecise, they are language and world dependent. To include such a set of variable phenomena into our core linguistic facts only seems to muddy the waters. If the goal is to determine the universal set of rules responsible for natural language, or the ontological attitude needed to describe them, then pragmatics can be safely set aside.

Let the above objections serve only to question the pluralist agenda. The issues, especially concerning the role of pragmatics, are controversial at best. My purpose is not to produce knockdown arguments here but rather to suggest how a scientist might abstract away from some linguistic material and arrive at a foundational project involving only syntax and semantics as "linguistics proper".¹² Consider Newton's theory of the tides. The earth's rotation and the gravitational pull of planets besides the earth, sun and moon are abstracted over. These factors can or do play a role in

modelling perspective I propose in section 3.1 and Part II aims to make this observation more precise.

¹¹These are technically obstruents which create small pockets of air and then release them in the production of loud consonants. Xhosa and Zulu are click languages of this sort.

¹²Of course, a different scientist could aim to abstract away from semantics by the same procedure. In fact, the arguments for 'autonomy of syntax' often presented themselves under this guise. See Sampson (2001) for an argument for the non-scientific/non-empirical nature of semantics based on similar reasoning.

tidal force (the ancillary force of gravity responsible for the tides) but they are not part of the core explanation of the phenomenon.

Moving on to the second line of argumentation for this view, let us consider the metaphysical pluralist claim. This argument is a response to an immediate objection along the lines of Postal (2003, 2009) as to the incompatibility of the various ontologies associated with mentalism, Platonism, physicalism and public language views. Stainton begins the pluralist apology in this way.

There is an obvious rebuttal on behalf of pluralism, namely that “the linguistic” is a complex phenomenon with parts that belong to distinct ontological categories. This shouldn’t surprise, since even “the mathematical” is like this: Two wholly physical dogs plus two other wholly physical dogs yields four dogs; there certainly is the mental operation of multiplying 26 by 84, the mental state of thinking about the square root of 7, and so on (2014: 5).

However, similarly to the result of thinking about the square root of 7, I do not believe this is analogy is rational. Mathematical reasoning does indeed involve mental operations, some physical examples (instantiations?) and the like but this is not usually how people conceive of “the mathematical”.¹³ On the standard picture, mathematics is considered an abstract science. Mathematicians study mathematical objects and rules which often outstrip the physical and the mental. The processes involved in mathematical thinking are certainly within the realm of psychology, but the fact that physical objects obey rules of arithmetic is not enough to hold arithmetic to physical characterisation. Does the question of how many dogs are in the union of an infinite set of dogs and another infinite set of dogs receive a physical interpretation? Stainton does acknowledge that natural languages display “interdependence” of these factors which perhaps “the mathematical” does not.

The main argument against incompatibility is that it rests on an equivocation of the terms “mental”, “abstract” and even “physical”. Once the equivocation is cleared up, it is argued, hybrid ontological objects are licensed. Let us briefly consider the different senses of these words proposed by Stainton. Physical₁ is something like an object under the purview of the hard sciences such as physics. Weeds, defined as unwanted plants, would apparently not count. On an extensional physical₂ definition, weeds show up since they have spatio-temporal and other physical properties. Mental₁ includes individual mental states such as pains and hallucinations. Mental₂ involves a specialised notion of secondary qualities conditioned by the mental but not identifiable with mental items. Aspects of taste and perception are suggested as examples of mental₂. Stainton uses the term “mentally conditioned” to capture this kind of mind-dependence. Lastly, he contrasts abstract objects *qua* Platonic objects, with what he calls “abstractish” objects, neither in the mind nor concrete particulars. Musical scores, models of cars and legislation form part of this latter category.

¹³Unless we say much more about the connection between the mental and the mathematical in terms of either intuitionism or finitism etc. See Part III for a discussion of some of these issues.

The argument goes that appreciating the physical₂, mental₂ and abstractish nature of natural language will absolve worries about ontological inconsistency. Consider some other members of this category of objects.

Indeed, our world is replete with such hybrid objects: psychocultural kinds (e.g. dining room tables, footwear, bonfires, people, sport fishing [...]); intellectual artifacts (college diplomas, drivers' licenses, the Canadian dollar [...]); and institutions (MIT's Department of Linguistics and Philosophy, Disneyworld [...]) (Stainton, 2014: 6).

I agree that such objects exist in all of the senses which Stainton describes but I think that there is something missing from the picture, something which breaks down the analogy between natural languages and their abstractish cousins. Natural languages are characterised as rule-governed by most linguists. These rules might be inner mental or mental₂ in nature but they also constitute types as *per* Postal and Katz's views. Stainton agrees to this much, "it is types whose meaning is compositional and systematic" (2014: 4). But types of this kind seem to have more in common with mathematical theorems and objects than they do with ordinary abstractish ones. This is not the place to go into detail here (see section 3.3 below and Part III for a different approach). But suffice to say that one feature of the former and not the latter objects is that there are uninstantiated types. There are sentences which have never been uttered, and thus are not physical₂ and not obviously abstractish (again, see Part III for a way of capturing the observation that they are in some sense).¹⁴ It is misleading to define such a broad ontological category and throw linguistic objects in with ordinary (abstractish) objects, with either extensions in space and time or secondary qualities with such extensions, in this way. More needs to be said about metaphysical hybrids and their potential subcategories before making any methodological claims based on them.

Furthermore, once the restricted domain of *linguistics proper* is reinstated, and with it goes methodological pluralism, many of the compelling arguments for the metaphysical variety are diminished (such as those involving phonetics and phonology). There are many insights to be had on the pluralist account, many of which I hope to retain, however I offer a distinct approach to the methodological aspects of linguistics, one which is equally compatible with various views on its ontology. The task of the subsequent sections is to present and defend this view.

¹⁴If we follow conceptualism we would consider them to be mental₂ though. This claim too is incorporated into the ontology presented in Part III.

Chapter 3

Three Grades of Mathematical Involvement

So far, we have seen a gradation (perhaps descent) of mathematical involvement in linguistics or *mathematisation* as I have called it. In what follows, I hope to impose some order and argue for sober methodological reflection on these issues. I shall approach the mathematisation of linguistics *via* a strategy of identifying grades of mathematical involvement for the grammars of linguistic theory, following a similar strategy proffered for modality by Quine (1976). The hope is that this will provide compelling argument as to which grade offers the best home for linguistic theory while avoiding the pitfalls of both conceptualism and Platonism. With each grade a further methodological burden is introduced. Importantly, however, as we shall see, this progression does not follow the traditional landscape of the debate in theoretical linguistics. The first grade is not just another label for nominalism, grade two for mentalism and so on. Although, the motivations behind Platonism seem to find a good place within the bounds of the third grade, this level of involvement by no means entails the Platonistic positions of Katz or Postal, as I hope to show. One way to think about the grades is that they represent a cluster of theories each with more commitments (in terms of connections to the target system) than the last.

Within the first grade of involvement, the mathematics involved in grammar construction is merely a helpful aid and not structurally or ontologically committing to the target system. The first grade of involvement proposes a quasi-instrumentalist picture of linguistic methodology (although this is not necessarily the case, see Part II.7 for a structural realist interpretation of linguistic modelling). On this account, grammars are indeed scientific in nature (*à la* Chomsky) but more akin to models than theories (*contra* Chomsky). On the second grade, the mathematics is part of the linguistic target system itself. Grammars are representational devices and directly represented (cognised, known or embodied) by speakers of the language. They are theories of linguistic competence. This view goes beyond mere modelling, for various reasons I will show. For now though, it is enough to understand that this grade does

impose systematic structural constraints on the reality which the grammar describes, in addition to material preservation and correspondence to internal mechanisms. Lastly, much like the previous grade, the third grade also holds that grammars are theories. However, on this level of mathematisation, grammars are mathematical or formal theories on the level of set theory, arithmetic or universal algebra. In essence, the mathematics involved in various grammar formalisms are enough to establish the reality or existence of linguistic objects without further empirical consideration.¹

Divorcing Methodology from Ontology

Part of the reason behind the failure of previous accounts in the debate on the foundations of linguistics is an over-emphasis on ontology. The problem has been diagnosed elsewhere (George 1989, and more recently McDonald 2009). The idea is that what determines the separation or subject matter of a discipline is not necessarily the ontological status of its objects but rather its approach or methodology.

For instance, in the section 2.2, we saw an argument from the ontology of sentences to the claim that linguistics is a formal science. Nevertheless, a linguist could grant that sentences are not located in space or time, and are therefore abstract objects or types, without conceding the point that linguistics is mathematics. The reason for this is that the question of with which discipline linguistics is aligned is not necessarily a matter of ontology. It is at least in part a methodological question. In other words, how the abstract objects, i.e. sentences, are studied and employed determines the science, e.g. as abstract objects in themselves or as convenient theoretical entities for scientific study. "It is not clear that having abstracta in the domain of a science is sufficient to make a science formal and nonempirical" (McDonald, 2009: 294). Put in another way (following George 1989), astronomy involves attributing numbers to the planets, for instance there are nine planets in the solar system. In a sense, astronomy involves *abstracta* since numbers are abstract objects

¹Yablo (2013) offers a related account of the three grades of mathematical involvement for scientific explanation. Yablo's three grades are roughly and respectively defined as follows: on grade one, mathematics has a descriptive role (something like the first grade on my view but more limited), on grade two it has a structural role and on the third grade it has a substantive role. He attempts to capture the substantive role in terms of a modal notion of *extricability*. We can think of extricability in terms of logical subtraction.

Logical subtraction *sometimes* yields a well-defined remainder, surely. *Snow is cold and white - Snow is cold = Snow is white*, I assume. For a generalization to be lawlike is what remains of its being a law, when we bracket whether the generalization is true (Yablo, 2013: 1014).

Scarlet is said to be *perfectly inextricable* from red (there is no logical remainder), congruence (of triangles) is *perfectly extricable* from equality in size, and action is *somewhat extricable* from bodily movement (it is evaluable in some worlds in which bodily movement does not hold and unevaluable in others). My grades, however, involve mental representation and other issues which are particular to the case of linguistics. Although inextricability is a useful means of capturing the claim of grade three.

but it is not concerned with abstract objects *qua* abstract objects as in mathematics. Neither astronomers, nor linguistics (I will argue), are directly interested in the properties of abstract objects but instead their concerns lie with how abstract objects might be related to natural phenomena or best model such phenomena, which is an empirical question (like the number of planets). George (1996) goes on to state that mathematical entities (such as grammars) can be identified empirically. He states that “[i]f I want to know the trajectory of a particle, I am engaged in an empirical inquiry whose goal is to identify a particular function” (George, 1996: 300). Similarly, identifying grammars is an empirical matter. I think that this perspective is too narrow for a number of reasons, many of which will be discussed in Part II. For one thing, it suggests a bottom-up perspective on scientific and linguistic investigation. However, in many cases mathematical entities (or models) were already developed independently (such as Post canonical systems) and then applied to natural phenomena. Thus, the situation is more like using a particular function to model the empirical target of the trajectory of a particle or whatever. Nevertheless, George’s analysis does suggest a useful perspective on the applied nature of linguistics and the status of mathematical entities within this enterprise.

Linsky and Zalta (1995) take the above suggestion further in claiming that abstract objects are not only a convenience for scientific theorising but a necessity. Their notion of abstract object is divorced from the usual Platonic concept but the idea remains that most sciences require some sort of abstract level of interpretation, whether it is the relatively benign use of numbers for characterisation in astronomy (and everywhere else) or the idea of species or types of animals in biology, *abstracta* of some sort seem to be a conceptual necessity (we will return to Linsky and Zalta at the end of Part III).

The above suggests that ontology is not the best way of capturing the uniqueness of linguistics or the Platonist claim that linguistics is a formal science. The grades of involvement I discuss draw the lines on methodological grounds and not ontological ones. This is not to say that my distinctions do not have any ontological consequences. In fact, such commitments are part of my caution concerning the second and third grades of involvement. Firstly, I shall introduce an useful analogy for the debate to follow.

Finding your way in New York City

Let us imagine that some person, let’s call him Kagiso, is planning a trip to New York City. Kagiso is the kind of traveller who enjoys consuming as much knowledge about a city before arriving as possible. Let us imagine further that he is on a budget and would like to see as many places across the five boroughs as he can during his visit. Thus, he is particularly interested in getting around by means of the New York subway system. So, he directs all of his efforts to understanding it prior to his arrival.

Now, there are different aspects of the system which might be of interest to a traveller. Someone could want to know the average time it takes to get from one point in the system to another, which stations are closer to one another and which routes need to be taken to get to different stations. One could also be interested in what the subway trains are made of, how they achieve their maximum velocity, what kind of mechanisms are involved in the braking system etc. One might say that there are different grades of answers which Kagiso might aim for in his understanding of the transit system. There are some questions to which he might not be able to get any answers or may be of no interest to him.

A convenient way of viewing the grades of involvement are along the lines of the following train claims.

- $C1'$: Structures of the representational device (phone GPS, subway grid map, picture etc.) are isomorphic (or homomorphic) to the structures of the actual subway system.
- $C2'$: Structures of the representational device are ontologically committing or substance equivalent to the subway system, i.e. made of the same stuff.
- $C3'$: Structures of the representational device include the actual mechanisms involved in the workings of the subway system (trains, sliding doors, ticketing etc.).

Now in terms of our example, Kagiso could be asking a number of kinds of questions relating to the claims above. For instance, he could want to know the average time it takes to get from Central Park in Manhattan and 81st (the National History Museum) to Flushing Avenue in Brooklyn. In terms of $C1'$ this would require some sort of map of the subway grid with a scale corresponding to real distances. If he were interested in an answer along the lines of $C2'$, he might require a 3D model of some sort (ideally made of the same substances as the real subway) as can be found at the New York Transit Museum (incidentally located in Brooklyn).² Asking a question would then involve running a real world simulation and viewing the result. $C3'$ would require information about how the trains actually get from Manhattan to Flushing Avenue, i.e. the inner working of their electrical makeup corresponding to how they traverse the system. $C3'$ might also involve the exact routes and paths (e.g. orange line, green line and how they connect) needed to get to and from these places. In other words, $C3'$ requires a mapping between the distances of the stations and the workings of the trains and routes in getting from one to the other which could include some physics or engineering information.

There is of course a way of answering the question about the distances of this route and all the possible routes without recourse to any of the claims above. There

²See Weisberg (2013) for a discussion on how the large scale Bay Area Model of San Francisco assisted scientists in rejecting the proposal to build a dam in the Bay Area.

could be a graph which accurately represents the distances between all the stations in the NYC subway system without respecting their “actual” routes, directions or interconnections. This graph would not be a map in the normal sense, since you couldn’t use it to find your way from one station to another. You could use it to accurately know the distance from Central Park to Flushing Avenue or from any place in the five boroughs to any other (within the subway system). Similarly, if you wanted to know the average time from one station to the next. All that needs to be preserved in this graph are the relative spatial and temporal relations between the stations and that structure is multiply realisable. For instance, getting to and from our designated place in Manhattan to the place we specified in Brooklyn could (and does) involve changing lines. Our graphs would neglect this detail. Taking the idea even further, there could be a permutation of the transit system of New York such that if we map NYC stations to stations in, say, the Kiev metro system according to which we would preserve the distances between the stations and average time between them. Now consulting the graph specified (or the Kiev metro grid) will not give us an answer in terms of $C1'$, $C2'$ or $C3'$ but it will *indirectly* track the spatial and temporal information (or structures) in which Kagiso might be interested, in this case distances and times (abstracting away from delays, commuter congestion etc.). I hope to show in section 3.1 that grammars operate in this indirect manner and provide a wealth of information despite their indirectness.

Specifically, the grades, I wish to propose, can be characterised in terms of the following claims.

- C1 : Structures of the grammars are isomorphic (or homomorphic) to the structures of natural language (or the linguistic competence thereof), i.e. directly structure preserving.
- C2 : Structures of the grammars are ontologically committing or material preserving to/with the structures of natural language.
- C3 : Structures of the grammars track actual mechanisms involved in language processing and comprehension.

The first grade of involvement is noncommittal on all of the above claims (but also compatible with them). The second grade of involvement is committed to $C1$, $C2$ and $C3$ while the third is only committed to $C1$ and $C2$. I will show that all of these claims (and thus the grades which are committed to them) are problematic in the case of linguistics. I also hope to show that neutrality on these issues is a virtue.

Notice that being a grade one advocate for the NYC subway system can yield genuinely useful insights. Kagiso will do fine in planning his trip with using a graph or map which only represents the distances, average time, and perhaps directions (maybe just a graph containing vectors) whether or not this graph actually shares the spatial relations of the system or its material makeup. Nor was his graph arrived

at by magic. Such a representation could have been devised for various reasons, perhaps it involves a simpler representation than an isomorphic structure would do or the pattern it uses is more user-friendly pictorially, perhaps Kagiso knows the Kiev metro system better.

Going according to the grades of involvement tells a story about mathematisation in linguistics in terms of how its practitioners treat their grammars. In other words, the grades track methodological claims rather than ontological ones but in so doing shed light on how the ontological claims developed from mere physicalism to abstract Platonism.

3.1 The First Grade: Models and Linguistic Reality

There is a growing literature on the nature of scientific modelling in philosophy. Although the role of models and their connection to the scientific enterprise has been much less explored than various accounts of or against scientific realism or the demarcation problem etc., modelling should be of particular interest to the linguist who, I will argue, faces a similar task to the empirical scientist in attempting to account for a natural phenomenon fraught with complexity by means of smaller more tractable representations of it.

Simply put, the first grade of involvement places our linguistic grammars at the level of scientific models. Their core aim is to capture salient features of linguistic reality, not necessarily to represent it in its entirety. This is achieved by various abstractions and idealisations, one of which is the notion of a “generative grammar” for modelling linguistic creativity. But let’s not jump ahead. Perhaps unsurprisingly, this position is much less controversial within the computational linguistics literature. “It is clear that to the extent that linguistic theories, i.e. grammars, aim to capture human knowledge of language, these theories are formal models” (Tiede and Stout, 2010: 147).

When a computational linguist provides a stochastic model or attempts to represent the next word in a grammatical chain as a finite Markov process, she is not necessarily making a claim about how human beings actually parse expressions (i.e. the exact route and train between two stations in the subway system). The model could have other evaluative benefits, such as predictive capabilities or efficient parsing complexity. Similarly, the model could be implemented in machine translation or other natural language processing uses. Explanatory models too work for a number of reasons in ways that do not correspond to adherence to *C1*, *C2* or *C3* above. I attempt to show this property of linguistic models or grammars in the following subsections.

My argument is that not only do grammars genuinely share a number of properties with scientific models but also that the only way to maintain the conceptualist approach of ascribing both infinity and a physicalist ontology to natural language is by

accepting that grammars are formal models of a target system, in this case linguistic competence. However, as I mentioned before, this grade of involvement is technically compatible with both a nominalistic and Platonistic ontology, i.e. grammars could be modelling idealised linguistic tokens, i.e. the output of linguistic competence (Devitt 2006), or an abstract mind-independent linguistic reality (Katz 1981, Postal 2003). I will not take a stand on these issues here (see Part III for an ontological account).

What are models?

So, what are models? And how do they relate to reality? One place to begin is by appreciating how modelling differs from other types of scientific theorising. The basic idea is that a model is an *indirect* representation of a target system or some aspect thereof, in this case natural language. The model bears certain resemblance relations to the target system such that stipulations within the model reflect aspects of the target system. For Godfrey-Smith, “the modeler’s strategy is to gain understanding of a complex real-world system *via* an understanding of simpler, hypothetical system that resembles it in relevant respects” (2006: 726). My claim in this section is that a grammar is precisely this sort of device and therefore that linguists find themselves in the modeller’s position with relation to natural language. By designing grammars which generate or constrain the grammatical output of a given language, linguists create small hypothetical systems which reflect or resemble structural descriptions of that language *via* rules that comprise the grammar. In this way, models, or grammars in this case, are theoretical intermediaries.

In order to see how this works, we should appreciate that mathematical models are essentially abstractions. They are abstract objects. They are designed to simplify a target system which otherwise would be too complex to approach scientifically (i.e. precisely). They might have various aims, simplification is one, explanation might be another, prediction yet another. Sometimes these aims can come apart. In his classical treatment of scientific modelling, Giere (1988) held that models were idealised structures (or abstract objects) aimed at representation of the target system in the real world. As previously mentioned, these structures or model systems bear resemblance relations to the target system. We might be interested in structural relationships such as various morphisms to capture this resemblance as in $C1$ (although Giere preferred a less formal account of the relation).³

The above picture is related to the semantic view of scientific modelling (Suppes 1960, van Fraassen 1980).⁴ Here the ambiguity of the term “model” is utilised. In

³However, another way to think of what a model is involves an analogy with fictional worlds, pretenses or ways that the world *could* have been (Frigg 2010). This view breaks down the connection with model theory in mathematics. In this way models are akin to the fictional worlds of Sherlock Holmes or Luke Skywalker. Counterfactual analyses are also generally connected to the type of representation involved in modelling. For instance, Giere (1988) affirms that model systems are systems which would be concrete if they were in fact real.

⁴As opposed to the then popular “syntactic” accounts in which scientific theories were considered

mathematical logic, a model is a set-theoretic entity with a domain of elements (or universe) and a relation which holds between those elements. “A model, basically, is a set of objects (and relations between them) that functions as an *interpreting structure* for a set of sentences.” (Godfrey-Smith, 2006: 727). Suppes (1960) held that scientific models and the logician’s set-theoretic models were one in the same. It is, thus, possible to talk of “truth in a model” or under an interpretation. Our scientific theories are then interpreted according to models with certain relations and structures. The central concepts here are “truth” and “satisfiability” and thus we might aim for a truth-preserving correspondence with the target system. So far we might still be tempted to accept something like *C1*. If grammars as models aim for truthful representation, then the target system might have to be beholden to the structures posited by the grammar. I will show this not to be the case in the subsection below.

In the following discussion, however, we will follow the literature in conceiving of as models belonging to a heterogeneous class of objects which includes physical models (used in biology and chemistry), scale models (used in engineering), computational models (used in population studies and computer science more broadly) and mathematical models (used everywhere and especially linguistics). Mathematical models, of which grammars are a proper subset, can be conveniently conceived of as abstract objects (whether they are set-theoretic, logical theories or fictional worlds).

Multiple Models: Against *C1*

There are a few ways in which a model can respect *C1* or direct structure preservation. Assuming a structure for the target system, we could require there to be a strong morphism, such as an isomorphism or homomorphism, between the model and the target. But as Frigg (2010) notes “[i]n order to make sense of the notion that there is a morphism between a model system and its target we have to assume that the target exemplifies a particular structure” (254). In Universal Algebra, morphisms are mappings between the relations and elements of formal algebras or structures.⁵ Thus, on this reading of *C1* we are forced to attribute a structure to the target system, i.e. natural language. This, however, is an unavoidable prerequisite for the modelling process.

The choice between structure mappings is also not arbitrary. Establishing an isomorphism places a much stronger constraint on the relationship between a model and a target system than does a homomorphism or weak homomorphism, since an isomorphic relation requires a homomorphism and an inverse morphism.⁶ The point

to be consistent sets of sentences in formal languages (“theories” in the logical sense).

⁵An algebra is a pair $\langle A; F_i \rangle$ such that A is a nonempty set called the carrier of A or the universe of A and $F = \langle f_i : i \in I \rangle$ are the (indexed set of) basic operations on A which are functions defined on A .

⁶If f is an isomorphism, $f : A \rightarrow B$, then there is an inverse morphism f^{-1} on $f, g : B \rightarrow A$ such

is that some kinds of morphisms will place an added burden on the model in terms of its relationship to the target system. For instance, if a strong morphism is required (such as an isomorphism), then all of the elements of the target system will have to be interpreted into the target system in a structure preserving way. Compare the structural relationship posited between two models connected by bisimulation in modal logic. In this case, the two models simply make the same modal formulas true despite potentially diverging greatly in their internal makeup (the simplest case involves a model with a one world cycle and a model with two worlds accessible to each other).

But standardly, as we have discussed above, models involve abstractions and omissions from the target system.⁷ Of course, we could have a circumscribed domain or in our case a proper subset of natural language, such as syntax, as a target. In this case we would need to establish an isomorphism between the model structures and the syntactic structures of various natural language constructions. Again, the choice of mappings becomes important.

The considerations above point to conceiving of $C1$ as introducing a range of relationships between the model and the target system. On one side of the range, are the strong structure mappings and on the other rather weak equivalences. In all of the cases, the theorist or modeller is attempting to establish a correspondence between the structures of their models and the structures of the target system.

The reason $C1$ often fails for models is due to one of the most common properties of models, they are multiply realisable. This property goes in both directions. Different models with different structures can be used to model the same phenomena and the same model can be used to model different phenomena. In terms of the latter scenario, consider a mathematical model of a pendulum and an identical model of a certain circuit.

The mathematical structures view seems committed to identifying both the pendulum model and the model of the electrical circuit with the mathematical structure they have in common and, thus, to insisting that the pendulum model and the model of the circuit are one and the same model (Thomson-Jones, 2012: 768).

Similarly, $C1$ identifies grammars by their structural/mathematical properties. But one could conceive of the same aspect of a grammar modelling two distinct natural language constructions (as in the pendulum and circuit case). Philosophers of science are wont to find additional means of model individuation. For Thomson-Jones (2012), models are constituted by sets of propositions. For Weisberg (2013),

that $f \cdot g = id_B$ and $g \cdot d = id_A$ where " id " is identity and " \cdot " denotes function composition.

⁷Hence the movement, in the philosophy of applied mathematics, to adopt partial morphisms which allow for undefined elements from the model to the world (see Part II.6 for discussion of these techniques in semantics).

the modeller's construal of the model differentiates between models with the same structures. Nevertheless, more than just structural equivalence is often needed to distinguish between models.

More pertinent to the case of linguistics is the scenario in which the same phenomenon can be represented by multiple models with non-equivalent structures. This is the famous problem of equivalent grammar formalisms initially presented by Quine (1972). Take two weakly equivalent grammars, phrase-structure grammar and tree-substitution grammar, for instance. These two grammars generate the same sets of sentences or in Quine's terms are behaviourally equivalent. The problem is that since they are empirical adequate or generate the same sentences, there is no way of deciding which grammar is the correct description of the target (in Quine's critique, the target would be mental states of language users). We will return to this problem in section 3.2 and a solution will be proposed in Part III. For now, it serves to question the need for a model or a grammar to respect anything as strong as the constraint on structure preservation exemplified by *C1*. This is also not a merely theoretical worry. Recently, there have been a flurry of formal proofs of weak or expressive equivalence of various syntactic formalisms such as tree adjoining grammar (Joshi), generalized phrase structure (Pollard) and categorial grammar (Steedman). Furthermore, Chomskyan syntax can also be shown to be equivalent to these formalisms (see Michaelis, 2001; Mönnich, 2007).⁸

There is another related worry stemming from the modelling literature. In Weisberg (2007b), he discusses a particular kind of modelling strategy, namely *multiple models idealisation*. This practice involves constructing many connected but incompatible models each of which focuses on one or more aspect of the target phenomenon. This strategy differs from other kinds of idealisation "in not expecting a single best model to be generated" (Weisberg, 2007b: 646). Naturally, structure preservation or one-to-one correspondences are not appropriate within this practice. Since scientific theories have diverse goals such as accuracy, simplicity, predictive power etc. and the construction of one model to fit all of these criteria necessarily involves "tradeoffs", this approach offers the theorist a way of meeting all of these objectives separately. "If a theorist wants to achieve high degrees of generality, accuracy, precision, and simplicity, she will need to construct multiple models" (Weisberg, 2007b: 647). This practice is common in climatology, ecology, biology and population studies.

If we consider the various models used in the service of linguistic theory and aimed at natural language, this picture seems to further militate against strict adherence to *C1*. Pragmatic models such as Stalnaker's model of common-ground and the conversational context or Lewis' scorekeeping in a language game are distinct from optimality theoretic formalisms of phonology and generative grammars for syntax. Even if we stick to the *linguistics proper* abstraction, generative enumerative syntax

⁸Or rather Stabler's (1997) formalisation and interpretation of minimalist syntax has been shown to be equivalent to some of the above grammars.

(based on proof-theory) and model-theoretic semantics (based on model-theory) embody distinct mathematical properties and formalisms, yet they both serve to capture an interconnected part of linguistic reality. Within syntax alone, we could conceive of the nested adequacy conditions of Chomsky (1965), mentioned in section 2.1, as multiple models with distinct goals. Although these models might not be strictly incompatible, there can be no notion of a structure preserving mapping (between model and target) when multiple models with distinct structures are being used to model a single target.

Thus, the idea that models or grammars have to preserve the (assumed) structure of the target system through mappings or morphisms is too strict and not a necessary condition on linguistic modelling.

Idealisation: Against $C2$

In this section, we will see how the rejection of $C2$ or material preservation can rescue the conceptualist position from inconsistency. Let us return to the modelling process *à la* Weisberg.

Modeling [...] is the indirect theoretical investigation of a real world phenomenon using a model. This happens in three stages. In the first stage, a theorist constructs a model. In the second, she analyzes, refines, and further articulates the properties and dynamics of the model. Finally, in the third stage, she assesses the relationship between the model and the world if such an assessment is appropriate. If the model is sufficiently similar to the world, then the analysis of the model is also, indirectly, an analysis of the properties of the real-world phenomenon (2007a: 209).

At stage three, however, scientists often do not associate similarity of their models and the target system with “truth” or even approximately true descriptions. For instance, in biology, Fisher described a model of fictitious three-sex organisms to explain the emergence of two-sex organisms involved in sexual reproduction (Weisberg, 2007b: 223). In physics, Boyle’s law is usually explained by assuming that gas molecules do not collide. This is not strictly true since low-pressure gases do collide. They do not tend to reflect the collisions in their behaviour with the result that these collisions are not admitted into the model (see Strevens (2007) for discussion). And in linguistics, we are asked to consider an idealised linguistic community of speakers and hearers who know their language perfectly and are never error-prone (Chomsky 1965).

The tools of models are abstractions and idealisations of various sorts. These terms are sometimes used interchangeably in the literature (we will choose a side in Part II.4). The basic idea behind these techniques is simplicity or tractability. Abstraction generally involves the removal of extraneous or superfluous material of the

target system in the model. Idealisation sometimes additionally involves distortions of the real world such as the cases mentioned above.

The details of these strategies are not of particular importance at this stage.⁹ However, there are some important features of idealisation which are relevant to the discussion at hand. For one thing, it is not ontologically committing as *per C2*, especially in the non-alethic cases (strictly false models). Fisher's three-sex model does not commit the target system to the existence of three-sex organisms, nor does Boyle's law commit the physical world to non-colliding gas molecules and importantly nor does Chomsky's model commit us to idealised speakers. Even more simply, a physical model of DNA does not commit the structure of DNA to the specific dimensions or the material composition of the model, e.g. a metre in size and plastic or styrofoam in composition. Secondly, as previously mentioned, models need not be directly related to the target system under study. In economics, von Thünen proposed a model of an isolated state on fertile land cut-off from all communication and contact with the outside world by a barren wilderness around its borders. Mäki claims that the false assumptions underlying this model serve the purpose of "neutralizing a number of causally relevant factors by eliminating them or their efficacy" (2011: 50). Thus, actually true and causally relevant aspects of a real economy are removed from the model.

In a similar fashion, grammars understood as models do not commit us to any specific ontology of natural languages. In the syntactic models of FLT, we treat sentences as sets of (uninterpreted) strings. This idealisation in no way commits us to natural language expressions being strings or sets of strings. In the semantics literature, meanings are designated by functional types. Meanings are not mathematical objects or functions in reality. Tiede and Stout (2010) go further in claiming that we are not committed to natural languages being discretely infinite just because recursion is a feature of our generative grammars.¹⁰ Discrete or denumerable infinity is assumed or a "modelling choice" on their view (and mine). The features or properties of the target system which we want to represent in our grammars are productivity, systematicity and conciseness. The first two properties are familiar from the literature on compositionality (presumably they would hold that this principle too is a modelling choice). The last property is assumed to prevent overgeneration of grammatical expressions.¹¹

Now it is clear to see how the conceptualist can have his cake and eat it too.

⁹For some details with relation to linguistics, see Part II and Nefdt (2016).

¹⁰This is true for technical reasons as well, as recursion does not guarantee discrete/denumerable infinity.

¹¹Here they borrow from Savitch (1993) who shows why we might assume that languages are (essentially) infinite despite having no evidence for them not being simply largely finite. Savitch's paper is a formal attempt at capturing parsimony judgements in grammars, i.e. we treat finite sets as *essentially* infinite if this allows us to get simpler descriptions than we would if we treated them as finite.

The “core” property of linguistic creativity is modelled as discrete infinity, the latter being the element of our models of natural language competence. Here we can freely employ idealisations such as sets of uninterpreted strings as sentences governed by recursive rules. The target system, however, is not committed to the ontology of the model, i.e. sets and functions etc. Therefore, the target system can still be a physical object or brain-states of individual language users.

Postal (2009) addresses (and dismisses) this (or a related) possibility briefly. He claims that to understand “infinite generation” or recursion as idealisations of some sort is to illegitimately equivocate on the terms ‘idealisation’ and ‘recursive’. As opposed to the idealisation of say a frictionless plane in physics, this idealisation is more close to “one which claims the solar system has an infinity of planets” (2009: 110). Postal deems such idealisations “silly”. However, I would argue that if such an idealisation were useful to a physicist or astronomer or helped understand some other property of the solar system, then it would be a perfectly acceptable aspect of a model (statistical cosmology is full of such idealisation). See footnote 11 for a reference and suggestion as to why an infinity assumption could be a simplifying tool for a linguist even if the target system is in fact finite. In addition, physics abounds with such idealisation. See the example of the thermodynamic limit in section 1.

If we maintain a separation between models and linguistic reality, the former being capable of the abstraction required for infinity statements and the latter being capable of physical description, then the conceptualist movement can be rescued from alleged “incoherence”. See the figure below for an illustration of the current picture of linguistic modelling (adapted from the general scientific picture in Giere 1988: 83). There are two salient relations here, the first is between the formal description of the grammar and the grammar as a model itself, i.e. recursive phrase-structure rules or constraints on feature structures etc., and the second is between the grammar and the target system which it needs to resemble in some way (the dashed line indicates that this resemblance relation is intentionally left vague).

According to the above diagram, the formal descriptions of the grammar might involve things like MERGE or recursion thus committing the grammar to discrete infinity (or the capacity for such cardinality) but the target system is in no similar way committed. For example, in discussing the question of the size or cardinality of natural language(s), Langendoen (2010) claims that “from the fact that one’s grammatical model is closed under such an operation [iterative or recursive operations], it does not follow that the language it models is” (2010: 140). This would require an additional argument or proof.¹²

Langendoen’s claim corresponds to the picture above in which the grammar aims to resemble or model certain aspects of the target system, such as our ability to process and produce previously unheard utterances or the fact that there seems to be non-arbitrary cut-off point for creating distinct expressions *via* repeated uses of

¹²Langendoen (2010) does, however, go on to attempt to offer such an argument.

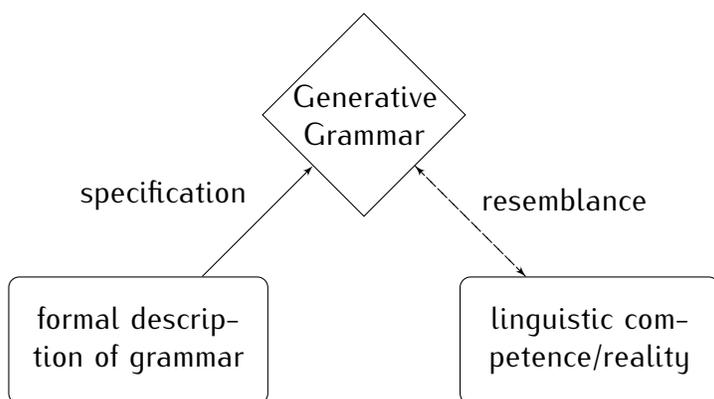


Figure 1.

conjunction or adverbial modification. Importantly, this feature of the model does not impose strict structural requirements on the target system as *per C1*. The resemblance relation could go in both directions, bottom-up or top-down, but in neither case does it commit the target system to formal features of the grammar in ontology as in *C2*. Hence, there is no incoherence here and this picture is derived from a similar view suggested for the rest of the sciences. In addition, if we return to the idea of “reverse-engineering” in the claim by Boeckx in section 1.3, we can appreciate the error of mathematisation in a new light. Attempting such a “reverse-engineering” assumes that the formal features of the model, such as recursion, requires an ontological interpretation in the target system such as an account of its evolution within the language faculty.

Grammars are abstract objects with an number of formal mathematical properties, these properties do not necessarily pertain to their linguistic targets in any ontologically significant way.

Lewisian Modelling: Against *C3*

Much like *C1*, a range is introduced by *C3* which starts from tracking specific linguistic mechanisms to more general cognitive mechanism involved with language (but also possibly other cognitive processes such as memory, movement, planning etc.). It is not clear, however, that a grammar needs to track any mental happenings whatsoever. The target of linguistic grammars could be outward linguistic behaviour or patterns which emerge from communities of speakers. In a sense, the individual idiolects (or I-languages) could determine the nature of these patterns, at least in part, but they could also be partly determined by external mechanisms such as linguistic conventions. Of course, different targets might in fact make different models desir-

able. For instance, on the Lewis' view, which will be discussed in this section, the explanation of subsentential elements of the grammar is of ancillary importance. A conceptualist might find such a view unappealing for this reason. In Part II.6. we will see how this view can be modified to account for such elements.

Before we discuss Lewis' view which has held sway amongst philosophers of language more so than theoretical linguists, let us briefly consider another example of how a model might not aim to track internal mechanisms responsible for production. A theorist might be interested in how people go about solving particular multiplication problems. There are two models which might be of interest here, standard "times table" multiplication or the more complex binary multiplication. There could be various reasons for preferring one model to the other such as respective processing times. The algorithm for binary is the same as that of standard times table or decimal multiplication but operates *via* three manoeuvres, namely $0 \times 0 = 0$, $1 \times 0 = 0$, and $1 \times 1 = 1$. So in binary, you merely replace symbols with other symbols with no carrying over as in decimal multiplication. The results are inter-translatable. In a sense then, it does not matter with which you choose to model arithmetic performances, since despite the different methods they turn out to be equivalent. In fact, mental multiplication might be "truly" captured by neither method. Nevertheless, the models could represent not only the results of such a mental calculation but other features such as timing correlations or common error explanations. Therefore, the actual mechanisms involved in multiplication might be quite different from the structures of the model and yet the model might indirectly correspond to those mechanisms nonetheless. In a sense, both methods exist at Marr's first level of the description of computational processes in which the function computed and the reasons why are stated. His second level states which algorithm actually computes the function (see footnote 17 for more details). We will consider positive views against interpreting grammars in terms of *C3* in section 3.2 but for now, we can simply appreciate that such a requirement seems unnecessary.

Furthermore, with relation to *C3* or the claim that grammars track actual mechanisms involved in natural language cognising, there is a compelling reason to think that this is not necessary. Lewis (1975) offers an account of how human beings use languages, construed as abstract objects (or functions from sentences to intensions), which is in the spirit of the first grade of involvement presented in this section. For Lewis, languages are abstract objects or functions which assign meanings to sets of strings (sentences). A language is then utilised by a community of speakers if and only if there is a convention in that community of *truthfulness* and *trust* in that language. The definitions of the terms are of no particular use to us here.

In terms of Fig. 1, the reality which is being modelled is linguistic communication. The formal description is given in terms of functions and sets of sentences (a grammar) and the resemblance relation is provided by a notion of "convention" which allows for a given formal object (or grammar) and not another to model linguistic communication accurately.

For an abstract language to be “realised” on this view, is for a community of speakers to possess a convention of some sort in that language. Sounds and sentences get their meanings relative to such a community of speakers and a meaning function (abstract language). However, there are a number of issues with taking an uncharitable reading of this proposal, i.e. taking it to be literally about the realisation of nonspatio-temporal abstract objects in the real physical world. Some issues are related to the specific six conditions placed on the notion of convention (see Gilbert (1989) for discussion). In terms of the dialectic of this chapter, one might wonder how exactly this “realisation” relation is to be construed on a literal reading in which abstract objects are considered to be outside of the causal nexus. In other words, how does a community of actual speakers access a function from sentences to intensions or *abstracta* in the necessary way?

On the modelling view, the above relation is unproblematic. On this account, languages are merely modelled as functions from sentences to intensions. Yalcin has a similar interpretation in mind with relation to formal semantics.

Semantic theory is not interested in the semantic value of properties of these abstract objects *qua* abstract objects. Rather, it is interested in an aspect of the question which of these abstract objects well-models what it is one knows, when one knows a language (2014: 36).

Yalcin goes on to remonstrate against Lewis’ impoverished notion of a language, especially its lack of the property of productivity (which he [Yalcin] considers to be a central desideratum of semantic theory). I think that the reason for this is that Lewis was rejecting *C3*-like reasoning and instead opting for an explanation of how a linguistic community uses or realises a public language (assigns meanings to its sentences). Therefore, the model was impoverished or rather simplified for this purpose and thus did not include additional properties such as productivity or systematicity and the usual compositionality facts. As previously mentioned, there is a way to retrieve these properties within Lewis’ general framework and I will discuss it in Part II.6.

A caveat. I am attempting to describe the grades of mathematical involvement not in terms of natural languages directly but rather in terms of the grammars of linguistic theory. Of course, various ways of specifying the nature and role of grammars could shed light on the ontology of natural languages themselves but as we saw, in this section, this need not be the case. So what does Lewis have to say about grammars?

In his earlier work on conventions, Lewis (1969) is non-committal about the specific nature of grammars. He states that they should somehow be finitely specifiable but infinitely capable, have a lexicon, and possess a generative as well as transformational component (so as to distinguish natural from formal languages which possess the latter in addition). He goes further in stating that grammars also need to assign interpretations to the set of generated grammatical sentences *via* compositional

projection rules. In Lewis (1980), he opts for a more broadly construed sentential account which importantly casts doubt on *C3*.

I use the word 'grammar' in a broad sense. Else I could have found little to say about our assigned topic. If it is to end by characterizing truth-in-English, a grammar must cover most of what has been called syntax, much of what has been called semantics, and even part of the miscellany that has been called pragmatics [...] You might insist that a good grammar should be suited to fit into psycholinguistic theory that goes beyond our common knowledge and explains the inner mechanisms that make our practice possible. There is nothing wrong in principle with this ambitious goal, but I doubt that it is worthwhile to pursue it in our present state of knowledge (81).

This is compatible with grammars being scientific models (or abstract objects in Giere's sense) aimed at representing only the class of linguistic objects or languages usable by various communities (and a few innocuous extras). The grammars do not seem to specify linguistic competence or *abstracta* in the traditional sense or "inner mechanisms" of language users, rather they specify functions that can describe or model the linguistic behaviour of language-using communities.

Lewis' position (as with my own) can be seen as a conciliatory intermediate position between the two different ontologies mentioned in the previous sections.¹³ In a sense, both positions, conceptualism and Platonism, are correct. They both tell a part of the story. Understanding the whole story involves appreciating how these ontologies connect with one another. If we appreciate a language as an abstract object in the sense of it modelling the patterns of speakers/hearers, we can connect the two ontologies. As Yalcin mentions, we are not interested in languages as abstract objects *qua* abstract objects but rather as formal tools for modelling linguistic behaviour. So in this sense, a language is a formal object picked out by a linguistic community by means of the finite rules of the grammar. It might be somewhat misleading to describe the view thusly. In the parlance of scientific modelling, a natural language, which is an abstract object for Lewis, is used to model the particular linguistic conventions of a given community. Substitute Lewis' *language* for grammar

¹³Despite strongly opposing Lewis' view in print, Chomsky's generative project can be characterised in a similar fashion to Lewis' project defined here. Consider Pullum's (1983: 449) remarks on generative grammar.

The discipline of generative grammar is founded on a crucial distinction between a language considered as a formally specified set of structurally described strings and a language considered as a behavioral repertoire. A grammar is by definition a membership specification for a language in the former sense. Chomsky has not provided any new definition of the notion "grammar" that severs it from its essential function of specifying the membership of a language. Therefore I continue, as is normal in generative work, to use the theoretical term "language" for the set of strings (or string/structure pairs) defined by a grammar, and not for anything ill-defined such as the set of dispositions toward verbal responses that characterizes a particular language user.

here and his *languages* for rule-governed linguistic behaviour or competence and you have the first grade of involvement for linguistics.¹⁴

By taking the scientific modelling route, as I have, we are also less bound by or susceptible to the specific follies of Lewis' account, in terms of specific conventions and their realisations in linguistic communities or languages as abstract objects in the strong ontological sense. Rather we are using models or grammars as tools for describing (predicting, reproducing etc.) representational systems which bear some resemblance to the physical reality of natural language, be it mental or social. The *Language* is given to us by the formal grammar which in turn sheds light on the *Languages* we use in our daily lives, in the construction of said grammar. My position allows for some neutrality on whether we are modelling mental processes or states or the behavioural output of public languages as in Lewis (1975). We might discover that the features of our models do track the internal mechanisms of individual language-users but this is by no means entailed by the modelling process.

I will close this section by offering two reasons for why the conceptualist and the Platonist might be reluctant to accept the first grade of involvement respectively. Firstly, as previously mentioned, the first grade does not entail any specific ontological position. It is perfectly compatible with our grammars modelling abstract objects themselves, theories (in the logical sense) or bee dances for that matter and not only the human mind/brain. Thus, it is a weaker claim than the conceptualist position in this sense. I think that this is an advantage of the view but it might escape an important aspect of the Chomskyan project, namely representationalism. I will return to this point in the next section.

Then, a Platonist might object that this story is all very well for generative grammars and denumerable infinity but the Vastness proof of Langendoen and Postal

¹⁴In fact, adherence to a strong interpretation of *C3* can lead us astray in some cases. There is a school of thought which takes infinity or fixed cardinality not only to be a modelling choice, as in the previous section, but to be a feature of the particular mathematical model used in linguistic theory, namely Post canonical production systems. Thus, infinity is an artefact of the model. It is obvious that not all the artefacts of models should receive interpretation in the target system. Especially if productivity facts can be captured by alternative formalisms which do not posit the putative property. For a specific example, Sampson (2001) criticises Chomsky's problematic "undue preoccupation with strings" in *The Logical Structure of Linguistic Theory* (1975). He points out that treating the syntax *via* derivations of strings and sets of strings is an unnecessary detour when phrase-structure grammars could be characterised with well-formedness conditions on trees directly. Furthermore, the derivational alternative forces certain untoward consequences.

Chomsky's approach forces him to impose two quite arbitrary restrictions on phrase-structure rules, namely, that no rule may rewrite any symbol *A* as either the null string, or as a sequence including *A*. Both of these forbidden types of rule frequently seem appropriate in describing real language, and under the alternative view of phrase-structure grammars there is no objection to them (Sampson, 2001: 156).

In this case, the model has features that the real world does not. In many other cases, the target has features which outstrip the models. Trying to find a home for every feature of the model as a mechanism or constraint might turn out to be deeply problematic.

shows that linguistic reality cannot be captured by such devices or models. So our grammars (defined generatively) if indeed they were models, as I have argued, would not be serving the purpose for which they are intended, that is representing the target system since they cannot even capture its totality. To the latter point, I argue that the story I provided is completely compatible with model-theoretic approaches to syntax (such as Generalised Phrase Structure Grammar, Head-Driven Phrase Structure Grammar, Sign-Based Construction Grammar etc.) which do not presume any upper-bound or fixed cardinality on the set (or “proper class”) of natural languages. In addition, our models might be modelling a denumerably infinite subset of natural language and perhaps have no need to go beyond generative capabilities (see footnote 11 for a related strategy). All of these options and more are open to a theorist within this grade of involvement.¹⁵ Not to mention, the science of linguistics is accounted for in a way which is consonant with the natural sciences while explaining the formal aspects of grammars such as infinity and recursion in terms compatible with naturalism.

3.2 The Second Grade: Representational Realism

On the second grade of mathematical involvement, the linguistic rules and their various posits (such as PRO, traces or copies depending on your generative persuasion etc.) of grammars are argued to have greater significance to the physical system represented than indirect representation or modelling of some sort. Linguists sometimes speak of the rules of a grammar being “internally represented” on this view. Chomsky (1986a: 243) describes a speaker “equipped with a grammar” as someone who “internalizes a system of rules”. On the basis of such an internalisation of a rule system R , the speaker’s linguistic behaviour can be explained or predicted by the structure of R (i.e. $C1$). To glean how this level of involvement is starkly different from the position described in the previous section, Pylyshyn is particularly illuminating.

[D]espite the uncertainties, none of us doubted that what was at stake in all such claims was nothing less than an empirical hypothesis about *how things really were inside the head of a human cognizer*. We knew that we were not speaking metaphorically nor were we in some abstract way describing the form of the data. (1991, 232)

We can see from this quote that the second grade of involvement, or “representational realism” as Pylyshyn calls it, is a much stronger claim than first grade. Grammars, on this view, are really like scientific theories and their posits are of the same nature as atoms and quarks are in physical theory, that is actual features of

¹⁵Part II will explore the fecundity of this idea more fully.

the physical system. If a grammar posits a mechanism of *wh*-movement or a recursive rule like adjectival modification (AP bar in X-bar theory), then these are features of a language user's actual brain-state when parsing these structures. This is in line with some form of *C3* set out at the beginning of this chapter. The claim that linguistics will eventually be subsumed by biology or neuroscience seems less vague on a second grade of involvement understanding.

It is important to pause here to consider the difference between theories and models again. In the previous section, I took models to be indirect representations of a target system. The mathematical structures involved in the building of models or grammars did not necessarily reflect any structural features of the target system in kind. For instance, positing recursive elements in the grammar only modelled iterative constructions in natural language indirectly, therefore I held that natural language did not need to be committed to recursion. In fact, the move within generative grammar from the recursive structures of Post-canonical systems to the single set-theoretic merge operation is evidence of the fact that recursion is an aspect of the models or grammars which can change without the target system changing (presumably language didn't change when linguists moved from the Extended Standard Theory to Minimalism). Scientific theories, on the other hand, represent the target system or natural world *directly*. In other words, scientific theories tell us what there is in the world. If grammars are scientific theories, then the structures and posits within them are claimed to be actual features of natural language. On this view, recursion and infinity are aspects of natural language competence and indeed conceptualists often speak this way (as in the many examples shown in section 1.3). Such an interpretation of the role of grammars naturally lends itself to analogies with mathematical cognition which presumably involves similar structures. For instance, if merge is an evolutionary mutation, it cannot merely be a formal aspect of a model. In order for this claim to even begin to make sense, it has to be assumed to be a claim about actual features of linguistic competence or reality posited by the grammar. Thus grammars preserve the very structures of linguistic reality and the second grade of involvement is committed to some version *C1* as well as *C3* (we will see how it is committed to *C2* in section 3.2 below).

Evans (1981) would find himself on this grade of involvement, in my view. In response to a criticism (initially levelled at Chomsky by Quine (1972)) that weakly equivalent grammars (mentioned in section 3.1, with different internal structures but equivalent behavioural output) pose a problem for representational realism, Evans offers a dispositional account of tacit (semantic) knowledge. Tacit knowledge for Evans is inferentially insulated ("not even potentially at the service of any other project of the agent" (1981: 339)). The important aspect of Evans' dispositional account is that it has an empirical, testable component. Two weakly equivalent grammars create distinct dispositions, ones which have distinct explanatory power.¹⁶ In other words,

¹⁶Evans argues that given two weakly equivalent systems, one containing axioms or primitives and

there are ways in which we can divine which grammars are “internally represented” by an agent on the basis of certain dispositions elicited by the structures or posits of the grammar.¹⁷ Accounts such as this one were part of the motivation behind early psycholinguistics. However, the failure of hypotheses such as the derivational theory of complexity or the claim that “the complexity of a sentence is measured by the number of grammatical rules employed in its derivation” (Fodor, Bever and Garrett, 1974: 319), showed that this project was not the sought-after evidence for representational realism. In fact, as Devitt (2006) notes, even when there is some positive evidence for posits of a grammar such as constituent structure (like in the famous “click location” experiment), this offers no proof of the psychological reality of the rules unless we presuppose the truth of the second grade involvement (representational realism or Devitt’s Representational Thesis (RT)). These experiments could show no more than that competence respects certain structural posits of our grammars (Devitt’s minimal position M). Fodor *et al.* themselves could not find any place for internalised grammar rules in actual parsing, “[t]here exists no suggestions about how a generative grammar might be concretely employed as a sentence recognizer in a psychologically plausible model” (1974: 75).

The interesting fact about psycholinguistic or physiological evidence is that it plays no real role in grammar construction. The primary data for the grammar on the second grade tend to be native speaker judgements (as discussed in 2.1. above). Since Fodor *et al.* (1974), there has been a lot of research conducted on psychological and biological effects and interactions in language production and comprehension. For instance, Cowart (1989b) discovered that familial handedness can have an effect on grammaticality judgements involving subadjacency. “[R]ight-handed speakers without left-handed relatives are more sensitive to subadjacency violations (rate them as less grammatical) than right-handers that have lefthanded relatives” (Keller, 1998: 7). If linguists were indeed in the business of developing grammars *qua* scientific

another containing composition rules and constituents, the former unlike the latter will be unable to predict the human speaker’s ability to understand previously unheard or novel sentences. However, the dispositional account notoriously suffers from philosophical rule-following problems. See Kripke (1982).

¹⁷Evans’ position has been developed by Peacocke (1986, 1989) and Davies (1987) to involve a level 1.5 between Marr’s first and second levels of the description of computational psychological processes. The first level is the *computational* one which specifies the goals and reasons for the computation. The next level is the *algorithmic* level in which the nature of the computation is specified, i.e. how it proceeds and is represented. The last level is the level of *implementation* which specifies the physical realisation of the computation. Marr himself held that the competence–performance divide mirrored the computational and algorithmic distinction (and that critics mistakenly take Chomsky’s internalised grammar rules to be at the algorithmic level). Peacocke claims that level 1.5 “states the information on which the algorithm draws” (1986: 101). This level allows us to be agnostic about the exact algorithm employed in natural language cognising while still offering a way of capturing the claim that grammar is concerned with individual psychology. For this reason, I think that this level belongs in the first grade of involvement which similarly avoids commitment to C3 and the range it introduces (between Marr’s levels 2 and 3).

theories of I-languages or brain-states, then such data would surely be relevant to the task. Yet this and other types of physiological evidence seems to have no place in the grammars of the second grade.¹⁸ This is related to an argument presented in Soames (1984) to the effect that linguistics and psychology are empirically divergent or they require different sets of evidence for confirmation of their theories.

The problem is that representational realism (strong *C3*), and the second grade of mathematical involvement generally, confuses the nature of a grammar. In addition, this confusion has led to undue mathematisation of the object of linguistic inquiry. Grammars are abstract objects, mathematical entities or models as I argued in the previous section. These tools are distinct from what George (1989) calls “psycho-grammars”, or “the mental state of knowing a grammar” (91). By appreciating this distinction, one can appreciate that grammars tell us nothing about mental processes or how those processes are implemented in a physical system like the human brain (the job of “physiogrammars”).

It is often claimed that being in such a mental state involves mentally representing the grammar. This claim is compatible with the grammar being silent as to the nature of these mental representations. Indeed, the grammar is not itself a characterization of a system of mental representation; it is the object of a speaker’s knowledge, not a description of how that object is represented by the speaker (if it is) (George, 1989: 91).

Even if we do accept this grade of involvement and its representational realism, we are still left in some confusion as to the biological underpinings of the movement or in George’s terms how we have any traction on the physiogrammar. This is the Postal problem of how the features of the grammars which involve sets, sentence types and discrete infinity are supposed to be captured by a physical biological system like a brain-state. In section 3.1. I offered a coherent picture of this relationship in terms of models and their idealisations. I also suggested that Chomskyans would be reluctant to accept this picture. The reason for this reluctance is that I believe that they have a stronger structural connection in mind, a combination of *C1* and *C3*, as evinced by their adherence to representational realism, which takes the rules and posits of the grammars to be structurally committing and thus actual structures of the mind/brain of language users used during processing. In order to capture the structural and representational nature of the second grade, the more apt analogy seems then to be something along the lines of a measurement-theoretic account.

Grade 2.5

Measure theory in mathematics is the study of how numbers are systematically assigned to physical phenomena. The standard example is temperature. Numbers

¹⁸I thank Geoff Pullum for drawing my attention to this research and general point.

are assigned to measure the temperature of various physical systems/environments. Different measurement systems can be systematically mapped onto one another. In this case, Fahrenheit and Celsius metrics are related in a structure preserving way. More technically, in analysis, a measure is a function that assigns a non-negative real number to each of the subsets of a certain set. The function assigns 0 to the empty set and is additive (larger subsets are composed of smaller ones). For example, in algebraic topology, the Euler characteristic is a number assigned to the structure of a topological space that is invariant under various ways that the space is bent or curved. Once more, the Euler number is assigned to objects in systematic way, so that, for instance, all Platonic solids have the same number (i.e. 2) as a sphere (because they can be reshaped into a sphere).

An approach with similar aims (but a distinct history) can be found in the measurement-theoretic accounts of cognitive psychology and the philosophy of mind.¹⁹ Johnson (2015) proffers precisely this sort of analysis for generative grammar (in the spirit of Suppes 1960 and specifically Matthews 2007). It must be noted that measure theory in mathematics (involving σ -algebras) is distinct from the representational measurement theory often discussed in psychology and the philosophy thereof. One aspect of this distinction is that in the latter types of accounts two theorems need to be proven, (1) the representation theorem and (2) the uniqueness theorem. (1) ensures that the model can be numerically represented in a way that preserves structure (a mapping or homomorphism from aspects of linguistic competence to sets in our case). While (2) specifies the unique set of numerical representations that satisfies the former theorem (the exact measure function). Johnson argues that the Quinian problem of extensionally equivalent grammars is actually a virtue in light of a measurement-theoretic analysis of generative grammar. In the same way that invariance plays a role in topology, the structures that are equivalent according to various grammar formalisms can be taken as invariant and anything outside of this is considered to be extraneous artefacts of the theories in question. This is the notion of ‘notational variants’ discussed in Chomsky (1972) and Chomsky (2000). Johnson describes it as “two theories (formal grammars, etc.) are notational variants iff they are empirically equivalent, in the sense that [...] ‘they do not differ in their empirical consequences’” (2015: 163).

He goes on to claim that this notion should play a more significant role in identifying symmetries in linguistic theory and thus the meaningful or real empirical content of various theories or grammars, i.e. distinguish between the empirically relevant and the ‘merely artifactual additional’ structure. Importantly for our purposes, the Postal problem is not legitimate in the context of measurement theory. It would be absurd to ask whether numbers were “inside” or part of temperature or thermometers, or length measurements of physical objects for that matter. When we measure temperature, we are not committed to the existence of numbers even if they are used

¹⁹For an historical overview, see Diez 1997a and 1997b.

in the service of its measurement. In addition, measurement theory could deliver us from commitment to $C2$, since the mathematical apparatus involved in grammars such as sets are not part of the target system in any ontologically committing way. In this view, measurement is a type of modelling.²⁰

If this analogy is indeed appropriate, then we can start to understand statements like the following (often misquoted) one made by Chomsky (in the context of a discussion about Goodman):

“We don’t have sets in our heads. So you have to know that when we develop a theory about our thinking, about our computation, internal processing and so on in terms of sets, that it’s going to have to be translated into some terms that are neurologically realizable [...] if we want a productive theory-constructive [effort], we’re going to have to relax our stringent criteria and accept things that we know don’t make any sense, and hope that some day somebody will make some sense of them - like sets” (Chomsky, 2012: 91).

A fully-fledged measurement theory for linguistics will supposedly show exactly how sets can be mapped uniquely and systematically onto the structures of natural languages as physical systems. The dual requirement of a representation and uniqueness theorem helps to distinguish between this grade and the first described in the previous section. The models or theories of the second grade of involvement are a special class of the general type and place restrictions such as accurate structure-preserving descriptions and morphisms on the relationship between the model and its target system.

One of Chomsky’s staunchest critics, Christina Behme, admits to the coherence of such accounts for a notion of mathematical modelling in physical systems. But she adds that such a story is not available for a Chomskyan concept of I-language since “there is currently no proposal providing a systematic correspondence between neurophysiological structures in the brain and the elements of the set-theoretic linguistic model” (Behme, 2015: 33). So the problem is that there is nothing resembling either the representation theorem or the uniqueness theorem for linguistics currently on offer. Hence Chomsky’s hopeful statement above. Even Johnson admits that realisation of measurement theory in linguistics is some way off but he holds that the first steps in terms of fragments of linguistic theory can be amenable to such analysis such as “merge” (see Johnson (2015) for details). A stronger sentiment is found in Soames (1984) when he argues for the “empirical divergence” of linguistic theory from cognitive psychology (and, one can assume, neurophysiology by similar reasoning).²¹ The claim is that the formal structures of grammars are not likely to be isomorphic to the internal structures of the mind/brain, nor should we expect them to be (here

²⁰See van den Bogaard (1999) for a development of this idea.

²¹He also argues for “conceptual distinctness” or the claim that linguistics and cognitive science “are concerned with different domains, make different claims and are established by different means” (Soames, 1984: 155).

conceptual distinctness comes in, see footnote 21). In fact, as we glean from Fodor *et al.* (1974) the rules of the grammars tend not to be reflected in psychological reality and from the Cowart case, they also tend not to involve physiological considerations. The grammar itself is agnostic with relation to the structures of the psychogrammars and their corresponding physiogrammars.

Furthermore, there is a question of what exactly is being “measured” in linguistics. In the hard sciences, there are obvious candidates for such measurements such as temperature, length, mass, spin etc. In linguistics it is not as obvious what is being measured, if anything. Johnson (2007: 384) claims that “[m]easurements in linguistics typically are not quantitative, but instead concern such things as the grammaticality or acceptability of a sentence, its sound and meaning, etc.” This might very well be the case but such measurements are not necessarily malleable in the same way as quantitative measurements are. For instance, the acceptability of a sentence is not a fixed feature, it varies empirically with each individual speaker. He goes on to say that “[t]hese measurements are used to reveal and explore various patterns that exist within various interestingly clustered sets of sentences”. But once again, the features involved in linguistics are not statically presented but heavily influenced by context. Measurement within a controlled experiment in physics might yield useful results which can be extrapolated outside of these conditions but such controlled experiments in linguistics with grammaticality judgements, for example, have led to serious challenges. The dynamic turn in syntax and semantics is testament to the failure of linguistic isolation or static conceptions. The grammaticality of a sentence can change or shift when context is added often resulting in the evaporation of anomalies even as strong as category mistakes. Consider the pair, (1)**I’m the salad* and (2) *No no he was the turkey club, I’m the salad*.

For Johnson, a linguist measures the linguistic properties of a sentence, for instance, when she categorises its subject *NP* as an agent and so on. This claim is somewhat out of sync with measurement theoretic analyses which generally deal with *magnitudes* or *quantities*. The latter are properties which involve gradation and comparison, such as the properties of *being tall* or *being fast*.²² Properties such as *being an NP* or *being an Agent* have no such quantities and thus cannot be compared or measured in the same ways. Lastly, even if a measurement theoretic account were viable, it would still involve only the psychogrammar and physiogrammar as *relata*. Assuming the abstract grammar rules (involving merge and other devices) would somehow map onto physical structures is committing the confusion mentioned above in the context of George (1989).

In many ways, a measurement-theoretic analysis induces more questions than it provides solutions for in generative linguistics. This makes it not only more problematic but also more complex in light of simpler accounts such as the one provided

²²The work on gradable adjectives is one case in which measurement theoretic modelling in the grade one sense of involvement has been quite fruitful.

in the previous section. Much more work needs to be done to render this a viable interpretation of the generative programme or its use of discrete mathematical tools. Thus, $C3$ is still not viable at this stage.

The Problem with Core Grammar

Lastly, the purported adherence to $C2$ of this grade of involvement needs to be addressed. In the section 3.1, we saw how assuming material preservation or substance equivalence between models and their targets is not necessary. Idealisations, false-models, even measurement models are all counterexamples to the claim that it is. Yet on the second grade of involvement, there is a position which seems to incorporate an aspect of $C2$. In section 1, within the broader context of linguistic mathematisation, I quoted Hinzen and Uriagereka as stating that “the processes described by generative grammar are functions computed over physical symbols by a Turing machine implemented in the human brain” (2006: 71).

In my view, statements of the above nature (frequent across the generative linguistics literature) are aimed at a specific concept within the second grade of involvement called “core grammar”. As Pullum (1983) describes the posit,

Chomsky does not assume that the grammars actually internalized by humans are (necessarily) defined as possible by universal grammar (UG). A basic “core” defined by UG is involved, but there is also a “marked periphery” of additional special constructions and exceptional cases that are learned on the basis of experience and not shaped in the same way by UG (448).

This proposal relies on the much maligned competence–performance distinction initially presented in Chomsky (1965). The idea is that competence is constituted by a generative grammar and indeed represented in the mind of the speaker. However, what psycholinguistic experiments are sometimes tracking are the heuristic and stochastic devices responsible for immediate parsing and performance needs (algorithms in George (1989)) hence the empirical divergence. One way to think of the relationship between the two systems is that the performance system checks itself on the competence grammar as it processes language in real-time. Think of the competence grammar as the generative grammar box situated somewhere in the mind, around it are the various quick-fire linguistic responses to external stimuli. The box is responsible for checking whether or not a given input is well-formed. When we produce and interpret sentences on the fly, we generally do not rely on this box. But upon reflection we often consult it to “check” whether or not a given string of words is in fact grammatical (hence the divergence between grammaticality and acceptability judgements). The idea stems from the notion of “core grammar” as opposed to “peripheral mechanisms” in early generative syntax. To lend some credence to this idea, at the beginning of *Aspects*, Chomsky seemed to base his idealisation of the true subject matter of linguistic theory on precisely this distinction (i.e. an idealised

speaker-hearer in a homogeneous speech community with perfect knowledge of her language and immune from the vagaries of memory limitations and the like). As Stabler notes,

The linguistic idealization is apparently grounded in the empirical assumption that the mechanisms responsible for determining how phrases are formed in human languages are relatively independent of those involved in determining memory limitations, mistakes, attention shifts, and so on (2011, 70).

This is all very well. But the concept of a “core grammar” has never been precisely laid out in the literature nor has its separation from peripheral mechanisms. As Pullum notes, “it is not clear whether the word “grammar” should be replaced by “core grammar” at the appropriate points. Nor is it clear to me what difference it would make” (1983: 449). If we are to follow a traditional competence–performance divide, we run into infamous problems. For one thing, it is not clear that the line between ideal competence and actual performance can be drawn as sharply as it is suggested here. Various aspects of performance have been shown to be highly systematic and context has been argued to have a significant effect on grammaticality or acceptability judgements (see Jackendoff 2002, Cann *et al.* 2012, Baggio *et al.* 2012). Furthermore, the assumption of domain specificity (which comes with the competence–performance divide) cannot go unchecked. In fact, some fMRI studies have challenged this classical architecture of the brain–areas associated with specific functions of language (see Binder *et al.* 1997 for one such study). In addition, even with this distinction in place, we still do not have much traction on the physiological or biological aspects of the grammar, i.e. why this competence constitutes an I–language or brain–state when neurophysiological data seems to be irrelevant to its description (e.g. the objection raised above in the handed subadjacency case).

The problem for the second grade of mathematical involvement is that in the absence of any evidence of the psychological or neurophysiological correspondence required for a measurement–theoretic analysis (or something like it) or a precise notion of “core” mechanisms, the success of its grammars as scientific theories is unclear. Moreover, it conflates different conceptions of grammar such as the abstract grammar (what I identified with a model in the previous section) and the psychogrammar (involving mental representation). I tend to side with Higginbotham (1991: 559) in stating that at the current level of scientific knowledge in linguistics, the second grade of mathematical involvement of grammars is not *indefensible* but rather *inarticulate*, taking *C1*, *C2* and *C3* along with it. I go a step further, however, in claiming that it is therefore an unsound methodological position in light of better options, such as the first grade of involvement. I think that Katz (1981), Postal (2003, 2009) proceed in a similar fashion. The difference is that they offer something along the lines of what I shall call the third grade of mathematical involvement for grammars as a more sound footing for the foundations of linguistics. In the next section, I evaluate

this possibility and find it no more satisfying on methodological and philosophical grounds than the second grade but with an increased epistemological burden.

3.3 The Third Grade: Grammars as Mathematical Theories

So far we have been trying to characterise the nature of the linguistic enterprise according to the mathematical involvement of its grammars. Along the way we met with Postal's challenge (2003, 2009) of trying to reconcile (1) the physical biological aspect of natural language (or the empirical scientific status of linguistics) with (2) the formal aspect of its description. In 3.1. I argued that this can be done successfully. However, my view took no essential position on the ontology of natural language and thus could be compatible with something other than (1). For those who hold something like (1), the second grade of involvement (the topic of the previous section) was the next natural step. However, on that grade, I argued, it is not clear how to meet Postal's challenge among other things (at least as the view currently stands). There is still another option available to those interested in a coherent ontology for linguistics. This option takes the form of rejecting (1) outright and placing (2) at the forefront of the linguistic agenda.²³ So the third grade of mathematical involvement for grammars places linguistics at the level of a formal science. Importantly, however, I hope to show that this position is not exhaustively captured by the linguistic Platonism of section 3.2. and specifically does not necessarily entail its ontology.

A grammar, on this grade, is also viewed as a scientific theory but of a specific kind, namely a mathematical theory. The modelling picture of 3.1 (and Fig. 1) is thus truncated and the formal descriptions of the grammar, such as the proof systems familiar from mathematical logic literature along the lines of the Post canonical system of section 1.3, specify linguistic reality directly. Another way to put this is that the structural relation between the grammar and the target, as per *C1*, is identity. On this view, grammar construction involves intuiting or deducing aspects of an abstract linguistic reality in similar fashion to proof construction in logic or mathematics.

On this grade of involvement, in order to describe or explain certain (constitutive) properties of natural languages at the appropriate level of abstraction (types in lieu of tokens) such as recursion or infinity, mathematics is not only structurally necessary as in *C1* but also materially so as in *C2* since both the grammar and the target are equally abstract.

My argument against this construal of the relationship between mathematics and

²³This is not to say that proponents of this view deny the psychological or physiological aspects of natural languages. Rather they argue that *linguistics proper* does not concern such things. Those aspects are more relevant to psycholinguistics or neuroscience.

linguistics takes two forms. On the one hand, I will argue for a methodological distinction between linguistic and mathematical theorising. On the other hand, I will argue that even if this latter objection can be overcome, the third grade of mathematical involvement still does not entail linguistic Platonism, since interpreting linguistics as a formal science opens up myriad possibilities within the foundations of mathematics.

Methodological distinctness might seem obvious to some. It might be argued that it is clear that linguistics does not make similar claims to mathematics nor use similar methods to establish those claims, namely *a priori* methods. Linguistic grammars are concerned with natural languages, use empirical data and are thus scientific theories (or models) not formal ones. However, doing this might be begging the question against the Platonist or the adherent of a third grade of involvement story (not identical to the Platonist) as I have represented them here. In addition, assumptions about empirical correspondence can lead us astray as in the second grade examples of the previous section between linguistics and cognitive psychology. The Platonist claim is precisely that natural languages are abstract in the same sense as natural numbers are, linguistic claims are true of an objective (necessary) acausal reality and grammars are proof systems or mathematical theories describing properties and relations of this reality.²⁴

One reason for the methodological discrepancy between linguistics and mathematics could be that linguistics certainly seems to use mathematical tools in identifying the properties of its objects (as do many sciences) but it does not seem to mathematically define the objects of its inquiry or rather use mathematical methods. Or as I have argued in section 3.1, linguists are interested in grammars conceived of as abstract objects but not *qua* abstract objects. In mathematics, once you stipulate or prove the consistency or necessity of an object, its existence follows (consider restricted comprehension or ‘separation’ in ZFC set theory). In linguistics, simply finding a consistent set of rules is not enough. These rules have to model the structures of real-world languages or linguistic competence, i.e. contingent facts. The fields of mathematics and linguistics are thus methodologically distinct.

Of course, a theorist on this grade of involvement could accept all (or most) of this reasoning and still maintain that linguistics is a formal science of a slightly different order. Katz (1981) anticipates some objections similar to the ones I have raised above.

²⁴If linguistics and mathematics were truly methodologically indistinct, then we would expect certain questions in the philosophy of mathematical practice to be relevant to linguistics, but they are not. Azzouni (2005) discusses three explananda of mathematical practice that seem to have no analogue in linguistics. It is there argued that any account of mathematical proof needs to explain *conformity* (general agreement on results), *phenomenology* (“aha moments” in maths) and *conservative formalizability*, as Azzouni puts it, “the success of the *Principia Mathematica* program of Russell and Whitehead exhibited, among other things, the conservative formalizability of classical mathematics, as it then existed” (2005: 153). However, although all of these properties are questionable in the case of linguistics, it is unclear whether they are any less so in contemporary mathematical practice.

He holds that there is a “single faculty of intuition” responsible for competences of different *a priori* areas of knowledge and the varying abstract objects under their respective remits. As for the connection between linguistics and mathematics, he has the following to say.

It may be said, for example, that the practice of the grammarian and the mathematician are dissimilar in that the working mathematician, unlike the working grammarian, does not spend large amounts of time soliciting and collecting intuitions. Conversely, the grammarian does not make extensive use of formal deductive procedures (Katz, 1981: 215).

For Katz, this is all just a matter of degree and the comparatively short history of linguistics. Comparing linguistics as it is today (or in the 80's) to mathematics as it is today is like comparing logic in the time of Aristotle to contemporary mathematical logic. Eventually as the science progresses, we will rely on intuition gathering less frequently. It is interesting to see a parallel here with Chomsky's hope for the future in the previous section. As for the lack of deductive procedures, he argues, that this is misleading. If we are talking about the proofs within metatheory (about systems) such as soundness, completeness, incompleteness etc. then linguistics indeed does not involve too much mathematics of this kind (with the notable exception of the Vastness proof mentioned in 2.2.). If, however, we are talking about first-order theories or proofs then there is an analogue in linguistics. The derivations of our Post canonical systems or generative grammars are such devices and these are ubiquitous. Once again, he thinks that the metatheory will also come along as more formalisation (or rather mathematisation) occurs in the study of natural language (a prerequisite for metatheory in formal systems).

I think that this is an interesting idea, even if it is speculative at best. Certainly, the history of geometry have shown a progression from concern with physical spatial intuitions and with the rise of non-Euclidean geometry in the 19th century culminating in the Hilbert programme in the 20th, to an abstract science not essentially informed by real-world constraints. I am not sure, however, how to imagine a similar scenario with relation to natural language in which linguistic intuitions no longer play any definitive role in grammar construction.

What underlies the methodological distinctness claim in my view is another confusion in terms of the nature of grammars which relates back to George's (1989) characterisation. In the previous section, on the previous grade of involvement, grammars as models or abstract objects were confused with psychogrammars and the physiogrammars which underlie them. On this grade of involvement, grammars are confused with the abstract objects themselves. As mentioned in section 3.1, we are not interested in grammars as abstract objects *qua* abstract objects but rather as abstract objects *qua* models of linguistic phenomena.

This is my case for the methodological separation of mathematics and linguistics (I will pick this up again in Part III within a more ontological setting). We can now

move on to the second part of the argument, namely that even if we do accept the third grade of mathematical involvement for the grammars of natural language, this in itself does not entail linguistic Platonism as I have described it in 2.2.

The reason is that the third grade of involvement, the claim that linguistics is a formal science akin to logic or mathematics (but perhaps not identical), does not necessarily entail the existence of abstract mind-independent objects. The philosophy of mathematics offers many different approaches to the ontology and interpretation of mathematics including nominalism. It is not clear to me why a Platonistic linguistics is considered to be the default position for a coherent ontology or for an analogy with the formal sciences. In fact, there are a number of reasons for opting for an alternative picture.

For one thing, if we accept a Platonistic ontology for linguistic objects we face Benaceraff's famous dilemma. Benaceraff (1973) argued that there is a tension between the semantics and epistemology for any theory of mathematical truth. If we attempt to offer a standard account of its semantics (in terms of our best truth-conditional theory) then we have reference to abstract objects which moves us further away from a standard (causal) account of its epistemology.²⁵ An indepth discussion of this dilemma is not necessary at this stage and will be addressed in more depth in Part III. Suffice to say, that with Platonism comes an added epistemological burden. Not only is reference to abstract objects difficult to explain but knowledge of an acausal non-spatio-temporal realm beyond the physical is highly problematic.

This problem *inter alia* has prompted many philosophers of mathematics to opt for nominalistic accounts of mathematics, which do not posit abstract objects (Field 1980, Azzouni 2004). There are also structuralist accounts, modal (Hellman 1989) which only require *possibilia*, eliminative, which similarly to nominalism, do not posit mathematical objects, non-eliminative, which do but in a "places-as-objects" within structures notion of object (Resnik 1997, Shapiro 1997). And then there are varieties of Platonism, hard-line (Gödel 1944) and more light-weight versions (Linsky and Zalta 1995). This list is not exhaustive by any means.

For instance, we will consider a modified version of *ante rem* or non-eliminative structuralism (chiefly presented in Shapiro 1997) as a foundation for linguistic ontology in Part III.

So if linguistics is a formal science, as the third grade of involvement assumes it is, then there is no principled reason to opt for the naive Platonism of section 2.2. As mentioned at the start of this section, the third grade of mathematical involvement for grammars, in which grammars are mathematical theories, does not entail any specific ontology for linguistic objects (or their necessary existence). Furthermore, we still have to account for the empirical side of linguistic research and how abstract languages relate to everyday spoken languages.

²⁵It's no surprise that Katz (1995) felt the need to respond to this dilemma by arguing that mathematical objects needed neither a standard semantics nor a causal theory of knowledge.

3.4 Subconclusion

There are many advantages to carving up the positions on the foundations of linguistics and linguistic practice in the way that I have attempted to do in the previous sections. By appreciating the role of grammars on each grade of mathematical involvement, we can divorce methodological concerns from ontological ones. We can restate the conceptualist view on less “incoherent” or rather “inarticulate” grounds and resituate the debate outside the scope of metaphysical complications. I further argued that by understanding grammars as formal scientific models, we can resolve the objections posed by Katz, Postal and others of the linguistic Platonist persuasion as well as avoid talk of abstract objects in themselves. The grades also exposed and confronted the mathematisation of natural language by showing that with each grade additional claims (*C1-C3*) as to the significance of the mathematical apparatus were imposed.

In this opening part, I have attempted to resituate (and hopefully reenergise) the debate on the ontological foundations of linguistic theory as well as the scientific nature of the discipline. I have placed the role of grammars at the center of my three grades of mathematical involvement for linguistics by describing them in terms of three methodological commitments, *C1*, *C2* and *C3*. The purpose of the grades is to show that the traditional characterisations of the various positions on the foundational issues as well as contemporary practice need not forge a marriage between methodology and ontology, as these concepts can be shown to be clearly separable for linguistics. Furthermore, I argued that the first grade of involvement offers the linguist the path of least resistance, drawing from insights in computational linguistics and scientific modelling. I argued that the second and third grades are more problematic but certainly not beyond redemption or merit. Specifically, I considered and rejected a measurement-theoretic analogy for a methodological stance that takes representations as real posits of grammatical theories (on the second grade). Finally, I clarified the third grade of mathematical involvement by drawing from the philosophy of mathematics and the various possibilities it brings into the logical space of interpreting linguistics as a formal science. My hope is that, at the end of the day, whichever path a linguist chooses to take for the interpretation of her field, it will be marked more clearly by appreciating or at least considering some of the arguments presented above.

In the next two parts of the thesis, I will aim to provide more details of modelling in linguistics and an ontological framework which could support it.

Part II

Methodology in Syntax and Semantics

Chapter 4

Modelling in Syntax

The purpose of the following part is to explore the power of the modelling idea in linguistics. The generative tradition took the form of a scientific revolution in the middle of the 20th century. The techniques and methodology which came along with the movement claimed to place the study of language at the level of an empirical, naturalistic science which would eventually be subsumed by biology or neurophysiology. As we have seen in Part I, arguments have been proffered which challenged this claim on ontological grounds (Katz 1981, Carr 1990, Katz and Postal 1991), methodological grounds (Soames 1984, Hintikka 1999, Devitt 2006) and linguistic grounds from the various competing frameworks, some of which were spawned from the initial generative approach (Pustejovsky 1995, Prince and Smolensky 1993/2004, Kempson *et al.* 2001, Jackendoff 2002).

In this chapter, I offer a lens through which to appreciate the scientific contribution of the generative tradition in linguistics in terms of the first grade of involvement of the previous part. This account specifies the types of modelling practices that this framework brought to the study of natural language(s), namely minimalist models idealisation (Weisberg 2007), a type of modelling that is ubiquitous in the hard sciences such as physics and chemistry. I use the above claim to provide an explanation of how the diverse and competing approaches to linguistics, specifically of the dynamic variety (Cann, Kempson *et al.*), are related to the generative one and a continuation (as opposed to a revolution) of the modelling strategies of the initial scientific revolution in linguistics. I argue that the generative tradition can thus be appreciated for ushering this type of modelling practice into the study of language and more broadly construed in terms of it.

This analysis does not presuppose any evaluative benefits or disadvantages of specific modelling trends. In addition, it does not aim to exhaustively capture all the modelling strategies employed by linguists, only some of the salient ones.

In the first section, I discuss modelling in the sciences with a focus on the notion of minimalist idealisation in model-building. This is by no means an attempt at a comprehensive account of the vast and diverse philosophical terrain of scientific

modelling, of which I have no intention (or need) of chartering at this time. In the following section, I attempt to provide an analysis of linguistic modelling drawing from the core tenets of the generative programme from the initial Standard Theory (1965) to Minimalism (1995). I identify two types of idealisation, namely minimal generation and isolation, both of which I argue are species of minimalist idealisation. In the next chapter, I attempt to extend this analysis to the dynamic turn in syntax and other related frameworks such as Optimality Theory and the Parallel Architecture. Lastly, in section 5.3, I discuss frameworks or rather types of frameworks which do not build their linguistic models by means of minimalist idealisation.

4.1 Modelling and Idealisation

Scientific modelling is a burgeoning field within the philosophy of science. The idealisations and abstractions involved in modelling have been argued to be pervasive in the sciences and seem to inform and shape much theorising in fields from physics to biology (see van Fraassen 1980, Cartwright 1983, Suppe 1989). In this section, I will focus on idealisation as I believe it plays a central role within the modern linguistic approach to natural language.

The terms ‘idealisation’ and ‘abstraction’ are sometimes used interchangeably in the literature. I will follow Thomson-Jones (2005) in distinguishing between these concepts. Thus, idealisations involve misrepresentation of the target system or specific aspects of it, while abstractions merely omit certain factors. “[W]e should take idealization to require the assertion of a falsehood, and take abstraction to involve the omission of a truth” (Thomson-Jones, 2005: 175). Thomson-Jones cites Chomsky’s invocation of an ideal speaker-listener in the study of linguistic competence as a canonical case of idealisation.¹ Another case of idealisation is Fisher’s Principle in evolutionary biology that states that the sex ratio of most animal species is 1:1 based on a hypothetical model which postulates a fictitious three-sex organism.

At first glance, this definition of idealisation might seem at odds with standard semantic accounts of modelling, such as Giere (1988), which assume resemblance relations (often in the form of morphisms) between the model and the target system. However, the idea of resemblance relations still holds even in an extreme case such as the Fisher model, in the form of a hidden *ceteris paribus* clause. We assume that all other factors of the biological world are held constant for the distortion or idealisation to explain the evolutionary stability (or evolutionary stable strategy, ESS) of the 1:1 sex ratio. In this way it resembles a *reductio* or constructive proof

¹“Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its (the speech community’s) 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 this language in actual performance.” (Chomsky, 1965: 4).

in logic and mathematics, in which the laws of logic (such as noncontradiction) are held constant while an absurd hypothesis is entertained (and eventually rejected). We will return to the issue of explanation and *ceteris paribus* hedges briefly in the next section.

As previously mentioned, for Godfrey-Smith, “the modeler’s strategy is to gain understanding of a complex real-world system *via* an understanding of simpler, hypothetical system that resembles it in relevant respects” (2006: 726). The important phrase here is “in relevant respects”. The relevant features of the real world which the model resembles might not be the properties which we are aiming to explain directly, these could be distorted if the model resembles the target system in other respects. In fact, it is unclear how idealisation would operate if there were no resemblances at all between the models and reality. Imagine a distortion or idealisation inserted into a system which in no way resembles the real world or the laws of nature. Not only would it be extremely difficult to predict the effect of such a distortion but it would be unclear as to the role of its introduction in an otherwise distorted world.

Of course, idealisations may be introduced for a variety of reasons. Weisberg (2007: 641) identifies a common type of idealisation in the hard sciences called “Galilean idealisation” which introduces distortions for the sake of computational tractability. A frictionless plane in physics is often referenced as a case of such an idealisation. No such thing exists in the real world and yet the idealisation is extremely useful in theoretical and applied mechanics. Formal language theory in linguistics possesses similar idealisations. In this field, natural languages are taken to be sets of uninterpreted strings organised according to their complexity. Of course, no natural language is wholly uninterpreted, but this idealisation is essential for much of the work done in computational linguistics and the construction of various grammar formalisms. Before moving on to the nature of linguistic idealisations such as these, let us consider what role they might play in the explanation of linguistic phenomena.

4.2 How the Laws of Linguistics Might Lie

In the philosophy of physics, Cartwright (1983) famously argued that the explanatory power of the fundamental laws of physics lies in their falsehood. Her *simulacrum* account of explanation relies on the idea the fundamental laws are not strictly true of observable reality but only true of highly idealised objects of scientific models. Reference to these latter objects are usually prefaced with *ceteris paribus* clauses which impose conditions never actually fulfilled in the phenomenal world (or the world of appearances, surface form in linguistics). Intriguing though this idea might be, it is generally considered to be quite a contentious matter in the philosophy of science and physics (see Elgin and Sober (2002) for a contraposition of Cartwright’s claims).

Nevertheless, in the case of linguistics this account seems somewhat more applicable. As we saw in Part I, the conceptualism upon which the generative programme is based seeks ultimately to explain linguistic laws (or rules of the grammars) in terms of biological or neurobiological reality. Thus, linguistic models, which are constituted by grammar rules, are not true of real world languages (which rarely met the requirements of such rules) and it is not even clear how they could be true of actual neurobiological states (which involve neural processes and synaptic connections etc.). The explanatory power of linguistic theories lies in the rules of the grammars of idealised languages, or I-languages.

In this way, the rules of generative grammars can be characterised as one of Stainton's options for an explanation of the field in stating that "the practice is sloppy, loose talk –which is strictly speaking false, and will eventually have to be reconstructed as corresponding truths about mental states and processes" (Stainton, 2014: 8).²

Ignoring the pejorative connotations of the previous statement, the competence-performance distinction which rests on the idealisation of a perfect linguistic community, incapable of error, further suggests that this picture might not be inappropriate for the rules of generative grammar. Whether or not we adopt an additional idealisation of core grammar (see section 3.2 of Part I. for details) to which the rules apply or the faculty of language narrowly construed (à la Chomsky, Hauser and Fitch 2002), the rules or laws of linguistics are not true of surface expressions but rather of highly idealised and internalised linguistic structures of the grammars. In addition, generativists are insistent that the rules of the grammar do not pertain to expressions of public languages or E-languages but the I-languages which in turn stand proxy for mental states and eventually brain-states to be explained by neuroscience. They are similarly insistent that the requisite cognitive and neurological structural realisations are forthcoming. Thus, the laws of linguistics seem to be doubly mendacious, in firstly being directed at explaining idealised structures of idealised communities of cognisers and secondly suggesting as candidate targets of the models mere placeholders for later biological instantiation. In the following section, I will delve deeper into the nature of linguistic idealisation.

4.3 Minimalist Idealisation from ST to Minimalism

In this section, I investigate two kinds of idealisation both aimed at discovering the minimal causal basis responsible for a particular property or phenomenon of the target system. It is an idealisation in the sense I have been using, in that the models make no attempt to represent the target phenomenon in its complete state or "de-idealise" to include extracted phenomena. In other words, we misrepresent the target

²He eventually goes on to reject this interpretation of generative linguistics.

system as involving only the core causal factors we deem necessary for the explanation or generation of a given phenomenon or property. Weisberg (2007) describes minimalist idealisation as “the practice of constructing and studying theoretical models that include only the core causal factors which give rise to a phenomenon” or “put more explicitly, a minimalist model contains only those factors that *make a difference* to the occurrence and essential character of the phenomenon in question” (Weisberg, 2007: 642). If this were mere omission, then we would be able to reintroduce the abstracted phenomenon into the model.

Consider the frictionless plane example again in mechanics. If we include friction (or fluid/air resistance) into the model, the predictions will fail, since these forces result in a loss of energy and thus a loss in speed and acceleration among other things. Admittedly, these elements are reintroducible into the system (and perhaps indicative of Galilean idealisation). A better example of minimalist idealisation is presented in Weisberg (2013: 100).

A classic example of a minimalist model in the physical sciences is the one-dimensional Ising model. This simple model represents atoms, molecules, or other particles as points along a line and allows these points to be in one of two states. Originally, Ernst Ising developed this model to investigate the ferromagnetic properties of metals. It was further developed and extended to study many other phenomena of interest involving phase changes and critical phenomena.

I believe that the generative tradition was largely motivated by such modelling practices, specifically through two versions of minimalist modelling, which I call *minimal generation* and *isolation* respectively. I provide examples of each in the following subsections. Before doing so, however, I shall state (following Blutner 2011)³ five core characteristics of the generative tradition in linguistics.

1. **Autonomy of Syntax:** The idea that the core “generative” component in natural language production is the computational system which produces the set of grammatical expressions. This system operates independently of the semantic, pragmatic and phonological components of the grammar (or in Blutner’s words “there exists an encapsulated system of purely formal generalizations orthogonal to generalizations governing meaning or discourse” (2011: 27)).
2. **Universal Grammar:** The claim that despite surface differences between the world’s languages, there is a set of genetically endowed linguistic universals common to all possible human languages (developments such as the Principles and Parameters framework allow for external linguistic input to shape the initial settings of the grammar).

³Although I have substituted his third tenet for the Universal Grammar postulate and my description of the rule-based view is somewhat different to his.

3. Innateness Hypothesis: A rationalistic approach to natural language acquisition in which human infants are endowed with a linguistic system prior to encountering any input. Often motivated by the “poverty of stimulus” argument (for some interesting empirical support for innate linguistic biases in child language acquisition, see Culbertson and Adger 2014, Culbertson and Newport 2015).
4. Competence–performance distinction: Linguistic theory is concerned with a ideal linguistic competence and not necessarily with the various aspects of performance or actual parsing and processing in real-time.
5. Rule-based Representationalism: This is the view that the posits of the grammatical theory or rules of the grammar are actual features of the human agent or ‘cognizer’ (actual goings-on in her mind/brain) at some level of deep neuro-physiological embedding. To ‘know’ or have a language on this view is to have subconscious (tacit or implicit) access to these rules.

Blutner goes on to argue that a broad construal of the generative tradition in terms of these aspects (or similar ones) should encompass frameworks such as Jackendoff’s architecture of the language faculty, Pustejovsky’s generative lexicon and the optimality theory of Prince and Smolensky. Importantly for my purpose, the dynamic syntax framework rejects many (if not, all) of these tenets outright and therefore I believe that the extension of this approach under the auspices of the generative one cannot follow the same lines as Blutner proposes for the other frameworks. In other words, the dynamic tradition constitutes a genuine theory change. I do not intend to dispute this point or argue that dynamic syntax is generative syntax in disguise. Rather, I propose an even broader construal of the generative tradition, along scientific modelling lines. This analysis maintains the broad construal of Blutner’s proposal but extends it in terms of modelling strategies as opposed to theoretical posits, i.e. generative linguists and dynamic linguists (and linguists of the other generative persuasions) build their models in similar ways, using similar strategies. It is on to these strategies that the next section moves, while in section 5.3, I also mention contemporary frameworks which do not follow these practices.

4.3.1 Minimal Generation

Minimal generation is perhaps the most explicit version of minimalist modelling in linguistics. One criticism of the pre-Chomskyan linguistic paradigm (the Structuralism of Bloomfield, Hockett and others) was its alleged inability to explain linguistic creativity. By focusing on statistical or classificatory aspects of specific corpora (of actual speech), this approach limited itself to dealing with finite tokens of natural language and thus could not account for the linguistic creativity.

As we saw at the beginning of the thesis, it was in drawing inspiration from computability theory that the early generative tradition placed the notion of a generative grammar (subclass of a Post-canonical system or Turing machine) at the forefront of the discipline. The idea was to capture the discretely infinite set of expressions of natural language *via* finite means. In the Standard Theory (1965) or ST, phrase-structure rules performed this task. Certain rules allow for recursive structures such as $NP \rightarrow (\text{adj})^* N$. Think of this as a loop in a push-down automaton (the class of automata associated with Context-free or Phrase structure grammars) which allows for unbounded iteration and thus a discretely infinite set of new expressions, e.g. Thabo is intelligent; Thabo is *very* intelligent; Thabo is *very very* intelligent etc. The product of the phrase structure rules (or rewrite operations) contributes to the DEEP STRUCTURE or underlying syntactic form. This structure feeds into the TRANSFORMATIONAL component of the grammar which is responsible for surface forms of expressions (through movement and deletion). ST was a progression on the transformational grammar of Harris (1951) and Chomsky's *Syntactic Structures* (1957) with the addition of a lexical component or lexicon which received input from the phrase structure and inserted lexical items into the deep structure. Kernel sentences, simple declaratives devoid of any modification which can be combined to form complex sentences, were also eschewed in favour of deep structure which could represent surface forms more minimally. This framework affirms the fourth tenet of the generative tradition, namely the competence-performance distinction. "ST does not attempt to answer the questions of language perception and production [...] rather than directly relating meaning and expression, it relates them indirectly, through deep structure." (Langendoen, 1998: 242). A direct relation (or determining relation) of meaning to expression is a matter of a performance grammar on this view and thus outside of the scope of linguistic theory.

ST, I claim, involves an example of *minimal generation* in the rewrite rules of the phrase-structure. In order to explain the creative aspect of natural language (or the specific examples of iterative structures such as conjunction and adjectival modification), i.e. the property of interest, we postulate a finite set of rules which allow for recursive structure and thus for infinite expression. The rule formulation is not descriptive but rather generative in the sense that it is supposed to represent multiple structures and with recursive elements (indicated by Kleene star above) potentially infinite structures. There are other ways to account for creativity. For instance, we could argue that we have a very large set of stored linguistic expressions (wholesale, not piecemeal) and we instantiate these expressions when prompted by experience (see Evans 1981 for a comparison between the two methods). However, this explanation would not be a case of minimalist idealisation in the same way that assuming we possess a finite rule system capable of infinite output is. The definition I propose can be stated in this way:

Minimal Generation: The explanation of a complex linguistic phenomenon or prop-

erty is provided by a model which includes only the interaction of the smallest possible units underlying the phenomenon/property.⁴

The above suggests an explanatory dimension to the modelling of a given phenomenon *via* the least possible units. In fact, the syntacto-centrism (tenet 1) of the generative tradition can be understood in terms of minimal generation. The idea is that we understand ‘generation’ in linguistics to be a means of providing explanations to causal questions.⁵ For example, if we want to explain why language users are prone to judging certain kinds of sentences (displaying certain kinds of syntactic structure) as felicitous or not (grammatical), then we do so by stipulating the least amount of rules which generate that type of sentence to model this behaviour. Thus, the rules of the grammar/model facilitate knowledge of the real world system through idealised models. This is similar to explaining the superconducting properties of certain metals via the Bardeen-Cooper-Schrieffer model which involves phase transitions and a thermodynamic limit, i.e. the nonveridical postulation of infinite particles. Now it is an idealisation technique because this story is not always strictly true such that “the endeavor is *explanation*; the feature of idealization [...] is the deliberate falsification of the causal workings of the system whose behavior is to be explained” (Strevens, 2007: 1). Strevens describes how Boyle’s Law is usually accompanied by a “causally distorting explanation” which involves ignoring the long range attractive forces between molecules and the collisions they exhibit *inter alia*. Similar idealisation is involved in explaining linguistic phenomena on this account. For instance, garden path phenomena (as in the example below) are notoriously difficult to parse by real speakers, yet they do correspond to syntactic rules, as in the famous case below.

1. The horse raced past the barn fell.

The rules of the grammar do not strictly correspond to speaker judgements in these cases but rather following Cartwright they correspond to the idealised structure that is the speaker’s I-language or state of the language faculty. They are true of a model. Thus, the model/grammar is not a direct representation of the target system, since speakers tend not to be able to parse these sentences despite their grammaticality.

The Extended Standard Theory or EST of the 70’s (Chomsky 1973, Jackendoff 1977), introduced further minimalist idealisations into the generative approach. On this account, the phrase-structure rules are simplified even further to account for a broader range of linguistic universals (or phrasal categories) via the binary branching

⁴In a popular syntax textbook, Carnie describes the generative approach as “[t]he underlying thesis of generative grammar is that sentences are generated by a subconscious set of procedures (like computer programs) [...] The goal of syntactic theory is to model these procedures [...] These rules are thought to generate the sentences of a language, hence the name generative grammar.” (2013: 6).

⁵Of course, this is not to be confused with the “generation” of Post canonical systems which is concerned with generative enumeration of a set. This is not a causal notion.

of the X-bar theory. In contrast to the many phrase structure rules of transformational grammar and ST, we now have only three types of rules which generate all the requisite structures. The three rules are (1) a specifier, (2) an adjunct and (3) a complement rule, represented respectively below (where X' is a head-variable and XP, YP, ZP, WP are arbitrary phrasal categories determined by that head).

1. Specifier rule: $XP \rightarrow (YP)X'$ or $XP \rightarrow X'(YP)$
2. Adjunct rule: $X' \rightarrow X'(ZP)$ or $X' \rightarrow (ZP)X'$
3. Complement rule: $X' \rightarrow X(WP)$ or $X' \rightarrow (WP)X$

Now X-bar theory vastly overgenerates the grammatical or well-formed linguistic structures and needs to be reined in by various other devices (such as theta-grids etc.). It is an idealisation in the sense discussed above. Once again, the model of grammar homes on the minimal causal basis necessary for grammatical representation. In addition, we move closer to an account which respects the innate structure of the language faculty, the third core characteristic of the generative approach (mentioned in section 4.3). As previously mentioned, in *Aspects of a theory of Syntax* (1965), Chomsky differentiates between three nested kinds of adequacy conditions for a theory of grammar, each more inclusive than the last. The three related linguistic desiderata are (1) observable linguistic performance, (2) native speaker judgements and (3) language acquisition. The first is the class of *observationally adequate grammars* which are those grammars which only account for corpora or observed utterances of speech. Naturally, these do not give us much traction on (2) and (3). Chomsky then suggests a class of *descriptively adequate grammars* (DAGs) which aim to capture the psychological facts of native speaker intuitions, thereby addressing (1) and (2). However, these latter grammars are inadequate on count (3) and thus require us to ascend to the level of *explanatorily adequate grammars*. By minimising the set of rules which learners have to acquire, we approach the explanatory adequacy necessary for a story about language acquisition.

[L]inguistics was supposed to be embeddable into cognitive science more broadly. But if this is the case then there is a concern about the unchecked proliferation of rules—such rule systems might be descriptively adequate, but they would fail to account for how we acquire a language-specific grammar (Ludlow, 2011: 15).

The X-bar innovation also pulled in the direction of universality as the new grammar rules or schemata could represent a greater number of tree (or hierarchical) structures and thus capture more of the constituents of a greater number of world languages, again with minimal resources. As *per* the definition of minimalist idealisation, we are only interested in the core causal factors involved in grammatical production, i.e. the models of ST and EST only contain these factors. In the opposite direction,

Newmeyer (2002) describes the generative semantics project as attempting to model too much and his words are particularly illuminating within the scope of the current section.

The dynamic that led generative semantics to abandon explanation flowed irrevocably from its practice of regarding any speaker judgement and any fact about morpheme distribution as a *de facto* matter for grammatical analysis [...] Attributing the same theoretical weight to each and every fact about language had disastrous consequences (121).

Another way of putting this is that the models were moving from minimalist idealisations to more comprehensive representations of the target systems (often including pragmatic phenomena such as implicature). In the next section, I describe another variety of minimalist idealisation modelling, one which, I think, is crucially involved in both the Government and Binding (1981) and Minimalist (1995) approaches.

4.3.2 Isolation

Natural language, and the linguistics which attempts to study it, is a diverse object of inquiry. Any theory which aims at a comprehensive account of its nature has to acknowledge the diverse factors involved in its explanation. When discussing syntax, semantic considerations invariably enter into certain descriptions (often captured by selectional restrictions on lexical items), when doing semantics, phonological aspects can be relevant (e.g. prosody) or pragmatic features (implicature, context shifting, metaphor, sarcasm etc.). Standard generative grammar places syntax at the centre of the language faculty (state of human language competence in the brain) and banishes these other aspects into various post-computational spell-out or logical form. However, some proponents, such as Jackendoff in his *Parallel Architecture* (2002), jettison the syntactocentric account and describe the language faculty as involving multiple generative mechanisms and interface principles between them. Dynamic syntax too rejects the centrality of syntax but goes one step further than Jackendoff in rejecting its autonomy likewise. The models of the generative approach aim not only to identify the minimal properties which “generate” (in the causal sense of produce) the intended aspect of the target system but also the relevant causes involved in this generation. This is where isolation comes in. Isolation is the modelling strategy which involves isolating or separating out the specific types of causal explanations deemed relevant to the phenomenon we are interested in producing.

The scientific modelling involved in generative linguistics often includes a property known as “modularity”. Modularity is the property of a system which involves separating it into discrete, individual subsystems which contribute to the systems overall organisation and operation. *Isolation* is similarly the technique of building models of these separate subsystems independently of one another (or as much as

possible). One can think of it as the modelling technique which corresponds to the property of modularity.⁶ So the definition, I offer, is as follows:

Isolation: The separation of a system into distinct minimal causal models for the *generation* of separate (but potentially related) properties or families of properties.⁷

This type of idealisation not only involves compartmentalising causal explanations but also potentially neglecting certain relevant causal factors outside of a given module. For instance, in an economic model of national GDP, one could exclude the contribution of a particular industry or sector (say, the value of production in the textile industry) even if this industry does in fact contribute to overall GDP. Stabler (2011) describes the competence–performance idealisation of *Aspects* in a similar way. “That is, we aim to find domains with causal interactions that are relatively closed, domains that can be described relatively autonomously” (2011: 69). I argue that Government and Binding or GB (1981) can be described in terms of such a modelling strategy. In this theory, separate modules govern separate aspects of the syntax (and semantics). As before, the minimalist idealisations identify an even smaller set of properties (for maximum generality). For example, there are only three core levels of the grammar on this account, namely D-structure, S-Structure and Logical Form. S-structure is derived from D-structure and logical form in turn from S-structure. The latter derivation is governed by a single MOVE ALPHA transformation at both the D to S-structure level and the S-structure to LF level (as opposed to a vast number of separate movement operations in ST and EST).

Importantly, the GB framework distinguishes seven separate modules which govern or generate different aspects of the grammar, in line with the initial autonomy of syntax thesis (tenet 1 above). The phenomena in question might involve multiple modules interacting but are explained within their distinct causal modules (as in the hypothetical GDP case above). One important application of the government relation involves the notion of abstract case, such as nominative, accusative, dative and so on (considered to be a universal property common to all languages, although often unrealised in surface morphology). Governance (which is a relation between heads and their phrasal categories, involving the dominance relation of m-command)

⁶This is a somewhat more general account of “modularity” than is found in the canonical cognitive science literature, such as Fodor (1983) or Pylyshyn (1984). This is because modularity is posited as a genuine property of a system or set of systems. Hence the claims usually associated with it such as *domain specificity* and *inaccessibility*. Isolation, on the other hand, is an idealising technique used in the service of model-building.

⁷Mäki (2011) and Portides (2013) discuss isolation in models as well. Although their analyses involve conceptual omission or “screening off” of features of an actual system. They differ in that Mäki considers the isolation as a result while Portides considers it as a process within the model construction. In this way, my conception is closer to Portides’. I do not, however, include a conceptual act within my characterisation.

also interacts with theta-theory which encodes semantic and functional roles such as agent, patient etc. However, Case theory and theta-theory do not necessarily coincide, despite being related causal explanations for various phenomena. For example, in the latin sentence below, both the theta-grid of the verb 'to give' or *dare* and the case of the indirect object requires/selects for a dative noun form of *Brutus*.

(i) Caesar dedit pecuniam Bruto. (Caesar gave the money to Brutus)

In GB these explanations are independent of one another. The idealisations of the theta-theory do not include those of the case theory, or rather they offer orthogonal minimal causal structures to explain the occurrence of the indirect object 'Bruto'. In GB, Chomsky describes the overall grammar in the following way (which exemplifies isolation idealisation).

The system that is emerging is highly modular, in the sense that the full complexity of observed phenomena is traced to the interaction of partially independent subtheories, each with its own abstract structure (1981: 135).

Finally, the minimalism program or MP, as perhaps the name suggests, provides the most radical case of minimalist idealisation at work. MP is often described as a programme or approach as opposed to a distinct theory on the same level as GB or the Parallel Architecture.

Minimalism isn't itself a *theory* of the language faculty that as such would or could compete with other such theories. No matter one's theoretical persuasion, a *minimalist strategy* of linguistic explanation is something one can choose to be interested in or not (Hinzen, 2013: 95).

Thus, in many ways, MP is the canonical case of a modelling strategy as I have described it. In MP, we start our models with only what we "must take to be true" and then rebuild the system from this basis. Once again, we see the concept of *minimal generation* described in the previous section. In terms of isolation, MP maintains the generative approach's separation between form and function (or competence and performance). In other words, the structure of the language faculty is independent of its communicative role. Furthermore, the communicative or functional aspects of the grammar are isolated from the formal features which have an alternative causal basis and role within a theory of grammar.

Previously we discussed Chomsky's notions of descriptive and explanatory adequacy. In MP, a level 'beyond explanatory adequacy' (also called 'natural adequacy') is introduced. The goal of linguistic theory now becomes to explain language as a natural object (in the sense of being bound by the biological and physical universe, as opposed to the mathematical and conventional ones).

In principle, then, we can seek a level of explanation deeper than explanatory adequacy, asking not only what the properties of language are, but why they are that way (Chomsky, 2004a).

In MP, language is considered to be a perfect system, optimally designed since its grammar constitutes a perfect computational system *via* economy principles for syntax and semantics (economy of derivation and economy of representation, respectively). Lappin *et al.* (2000) argue that both perfection and optimality are unclear notions in this framework and should constitute serious challenges to MP's adoption by those linguists working within the GB framework. In terms of my dialectic, the difference between GB and MP is especially illuminating.

Throughout the modern history of generative grammar, the problem of determining the character of FL [faculty of language] has been approached "from top down" [as in GB framework]: How much must be attributed to UG to account for language acquisition? The MP seeks to approach the problem "from bottom up": How little can be attributed to UG while still accounting for the variety of I-languages [internalised language or specific state of the language faculty] attained, relying on third factor principles? The two approaches should, of course, converge, and should interact in the course of pursuing a common goal. (Chomsky, 2008b: 4).

Chomsky's distinction between "top down" and "bottom up" is not entirely clear. It can be, however, related to a topic in the theoretical physics and chemistry concerning what is referred to as "foundational" versus "phenomenological" approaches. The latter are the various frameworks such as GB, ST and the principles and parameters (P&P) which offer specific analyses of linguistic phenomena. Foundational approaches, on the other hand, aim to answer the questions concerning the reasons behind the use or application of a given formalism. This might involve the search for a set of first principles which independently motivate the use of certain theoretical tools or explanations.⁸ Hinzen (2000) offers a comparative analysis of the minimalist program and the principles and parameters framework along these lines. He states, among other things, that minimalism attempts to rationalise rather than describe the phenomena under study. Furthermore, it aims to discover general principles underlying explanations and avoid overly technical solutions. GB can be compared to MP similarly. Whereas the GB framework approached the constitution of the common linguistic substrate or Universal Grammar by asking 'how much' structure needs to be innate, MP asks the question of 'how little' structure is needed. The operation of merge (as well as *select* and *move*), which takes two items and creates a labelled set containing both of these, is supposed to be the minimal requirement on the productive capabilities of the language faculty. Our complex model of natural language syntax now only involves a single operation which serves as the minimal

⁸I thank Reinhard Blutner for drawing my attention to this point and research.

causal basis for the entire system isolated from other potential causal factors (such as functional roles, the conceptual system etc.). As we have seen, there are some interesting ramifications of the merge postulate, both evolutionary and ontological.

In this section, I have claimed that the generative programme in linguistics, from ST to MP, encompasses minimalist idealisation in the form of both minimal generation and isolation in the models of the various theories. I followed a Cartwrightian line in claiming that these techniques are indeed idealisations in terms of falsehoods not only because the rules of linguistic theory pertain to highly idealised models but also because these models are taken to stand as placeholders for the true descriptions of a future neuroscience. I now move onto extending this analysis beyond generative grammar (narrowly construed) to other frameworks and the dynamic turn in syntax.

Chapter 5

The Dynamic turn and other Frameworks

5.1 Other Generative Frameworks

Within the more broadly construed generative tradition in linguistics, we find many examples of both isolation and minimal generation, as I have described them above. Perhaps Jackendoff's parallel architecture (PA) serves as one of the best cases of both isolation and minimal generation and therefore a useful starting point.

One of the aims of Jackendoff (2002) is to better integrate linguistics within cognitive science. In order to achieve this aim, he rejects a number of components of the Chomskyan view of generative linguistics, for instance the syntactocentrism, or the view that syntax is the central generative element of language. Jackendoff holds that this was a mistake. In opposition to this view, he proffers a parallel architecture of the language faculty.

The alternative to be pursued here is that language comprises a number of independent combinatorial systems, which are aligned with each other by means of a collection of interface systems. Syntax is among the combinatorial systems, but far from the only one (Jackendoff, 2002: 111).

He goes on to describe each independent rule-bound and hierarchical system in isolation from one another. This analysis includes a reconceptualisation of semantics as "a combinatorial system independent of, and far richer than, syntactic structure" (Jackendoff, 2002: 123). Given this high level of modularity, we can glean a perfect case of isolation idealisation at work. Each system, phonological, syntactic and semantic are generated by independent structures. Due to this modelling strategy, the interfaces between these structures becomes of particular importance in terms of a holistic concept of natural language. For an idea of how this works, consider the concept of the well-formedness of a sentence. Within the frameworks of the previous

chapter, the syntax determined the well-formedness of a sentence and the other steps in the derivation (phonological and semantic) were somewhat epiphenomenal (recall Chomsky's famous *Colourless green ideas* example which was meant to show grammaticality outwith interpretability). In the parallel architecture, the situation is different. A sentence is only well-formed if it is so within each separate system and there is a well-formed interface between them.¹ Burton-Roberts and Poole (2006) take issue with this aspect of PA. They argue that trying to capture the structures of the modules independently results in a loss of the initial rationale behind those structures.

The term 'semantic' [...] is relational. It suggests that the module is distinct from the central conceptual system in being dedicated to specifying the SEMANTICS-OF something – expressions generated by the syntax, presumably. But this implies that those expressions have semantic as well as syntactic properties [...] Equally, a mechanism that specifies the semantics-OF syntactic expressions cannot be encapsulated with respect to syntax. Its rationale lies in syntax, being effectively 'interpretative' of it (as in models the PA claims to repudiate) (Burton-Roberts and Poole, 2006: 622).

The complaint is essentially that the models of PA neglect causally relevant material, i.e. part of what determines the semantic module is syntactic in nature or related to syntax. This, however, is consonant with my characterisation of isolation idealisation (in terms of falsehood). In this type of idealisation, causally relevant aspects are often ignored and false models are created for explanatory purposes. We will return to the interpretative nature of semantics in the chapter 6. For now, it suffices to appreciate the isolationist modelling of the PA, whether it can retrieve the connections with syntax (through interfaces) or not is not our chief concern here.

Nevertheless, despite the differences, the parallel architecture does maintain the autonomy of syntax (and phonology and semantics) as well as the UG hypothesis (although Jackendoff takes pains to divorce the concept from misinterpretations in section 4.2. of the book) and the competence-performance distinction (once again with some criticism of how the idealisation has "hardened" over the years).

Optimality theory (OT) is another approach in which minimalist idealisation is harnessed. Minimal generation is both an implicit and explicit device in OT. Explicitly, the formalism contains a GENERATOR which generates an infinite number of outputs or candidates for representation for each input of the grammar. The EVALUATOR component then chooses the optimal output from the set of outputs through a set of ranked, violable constraints or CON (in the sense that violations are permitted but those of higher level constraints count more than violations of lower level ones against the

¹Some of these interface principles or rules are constraint based, such as the head constraint (borrowed from HPSG) for the syntax-semantics interface or the required linear order mapping between phonology and syntax. There are also static idioms which bypass syntax entirely and occur between the phonological and semantic components.

potential optimal candidates). *CON* is considered to be universal (in line with tenet 2). In terms of generative grammar, it possesses an assumption “that there is a *language particular* ranking of constraints from a universal set of constraints” (Blutner, 2000: 2).

One reason for questioning the place of OT within generative linguistic modelling, as I have been describing it, is that it seems to be constraint-based as opposed to derivational or “generative” (in the proof-theoretic sense). It should be noted that GB also has a distinctive constraint-based flavour (more on model-theoretic syntax in section 5.3). However, importantly as Smolensky (2001) notes, “OT has been formulated in both derivational and non-derivational or ‘parallel’ forms. Both variants are coherent expressions of the theory”. The core idea in both cases can be explained in terms of minimal generation idealisation. The property of being an “optimal candidate” is generated directly by a universal set of inputs narrowed down by a minimal set of constraints. Blutner (2000) himself offers a bidirectional OT approach to semantics which is somewhat different to the generally unidirectional analysis of the generative tradition. Nevertheless, the key idea here is the minimal set of constraints. In OT phonology (where the framework received dominant status), the best analyses are the ones which generate a given typology of phonetic combinations *via* a minimal set of constraints and their relative rankings (see Hammond 1997, McCarthy 2003b). In OT, there is no room for extraneous constraints. In fact, the methodology is essentially concerned with defining the fewest and often most specific constraints necessary for generating optimal candidacy.

I think that this should be sufficient to display the pervasive nature of minimalist idealisation through both minimal generation and isolation within the broader generative tradition. It might be objected at this point that there is major theory continuity within the frameworks so far discussed and perhaps the modelling practices can be more easily explicable in these terms. I do not think that this is necessarily the case. Cartwright, Shomar and Suárez (1995) argue that theory and modelling are independent processes in the sciences. They argue that theories can serve as tools for models but are not to be defined by them (i.e. not just a collection of models in the van Fraassen (1980) fashion). Unfortunately discussing this version of instrumentalism will take us too far afield, although in a similar vein to Cartwright *et al.*, I will attempt to show, by example, that modelling practices can be held constant despite significant theory change. However, my method is the reverse of the one they take. While they argue that the London model of superconductivity underwent model change without theory change, I will argue that dynamic syntax utilises similar modelling strategies to the generative tradition while the theory has been shifted on almost all accounts.

5.2 Modelling Dynamics

In this section, I hope to extend my analysis of the modelling strategies employed within the generative programme to a rival approach, namely dynamic syntax (DS). In so doing, I also hope to provide an account of the theoretical differences between these frameworks and their genesis in terms of the choices of minimal structures within the models as opposed to a shift in modelling strategies *in toto*. As previously noted, I will not be disputing the claim that DS marks a significant departure from the theory presented in generative grammar. For instance, the competence–performance divide, representationalism and the autonomy of syntax are all unabashedly abandoned in this framework.

In an attempt to account for the deep context-dependence of natural languages, DS roots its idealisations in the “dynamics of real-time language activities” (Cann *et al.*, 2012: 359) where semantic factors inevitably affect any analysis. Thus, there is no autonomous syntactic component and the idealisations of formal language theory (which were heavily reliant on the alleged connection between formal and natural languages) are jettisoned in favour of a model of incremental semantic growth.

This theory does not characterise the surface (constituent) structure of a sentence, but instead models the process of assigning an interpretation to a string of words in a left to right fashion. In other words, taking information from words, pragmatic processes and general rules, the theory derives partial tree structures that represent the underspecified content of the string up to that point in the parse (Cann, 2001: 4).

For the point of illustration, let us return to the idea of well-formedness. We saw with generative syntax and the parallel architecture, two different but related notions of the well-formedness of a linguistic expression or sentence. In DS, the departure is more stark. Syntactic well-formedness is no longer the determining factor within the linguistic concept. For instance, in multi-person dialogues, certain sentences (or strings) can be plainly ungrammatical in isolation and yet give rise to well-formed structures.² Consider the example from Cann *et al.* (2012: 365) below:

Father: We’re going to Granny’s.
 Mother: to help her clean out her cupboards.
 Child: Can I stay at home?

²Although the notion of ‘string well-formedness’ is not part of this approach (i.e. there is no ‘membership problem’ in the Turing sense), the idea of a well-formed or ‘complete’ utterance is present (on the basis of which grammatical judgements are made). “We may take the concept of ‘complete utterance’ in some language *L* to be one for which it is possible to construct a propositional tree of type *t* from an uttered string of words using the lexical, computational and pragmatic actions licensed in *L* (Cann *et al.* 2005: 398).

Mother: By yourself? You wouldn't like that.

Various traditional locality requirements on pronouns or anaphors (such as *yourself*) are violated in this exchange and yet it is unproblematically interpretable and natural. Thus, the models of DS are built up from a basis that goes beyond the single-person and sentence level boundaries of the previous frameworks which we have discussed. In addition, the formalism acknowledges the word-by-word contribution within expressions and not only the final output of a derivational process (in this way following the path of unification-based grammar formalisms such as GPSG and HPSG etc.). "The way this is achieved is to begin from a goal associated with some very partial structure and progressively enrich that structure through the parse of a string of words" (Cann *et al* 2005: 33). Various techniques from semantics and dynamic semantics, such as underspecification and updates, are incorporated in order to accomplish this analysis.

The usage-based (parsing) elements of this formalism take it further from the abstract rule-based representationalism of the generative tradition toward a characterisation of linguistic knowledge as a type of "know-how". Thus, it seems as though the models of DS are vastly different from those of the generative tradition and indeed, in some respects, they are. However, in terms of the type of modelling employed by practitioners within this framework, I think some important continuity can be found. Primarily, I hope to show that DS does employ a minimalist idealisation approach to its models.

In order to see this, let us revisit the motivations behind some of the aspects of the theory change in DS. One of the leading motivations behind DS (and the dynamic turn in linguistics in general) is the apparent failure of static accounts to deal with phenomena such as ellipsis, anaphora and tense. In other words, the objection is that by focusing the minimal models on a static sentence and single person boundary, the generative tradition (and other approaches) has failed to capture the property of interest in these cases, i.e. acceptability judgements of speakers in many of the cases involving dialogue data etc. Furthermore, ignoring such data as performance error or dysfluency is claimed to result in incomplete models as well as only a partial approach to the language acquisition problem (or the 'explanatory adequacy' of the generative tradition) since young children are confronted with such data on a daily basis and it is systematic. "The effect [of ignoring the aforementioned data] will be that no single linguistic phenomenon will receive a complete characterisation" (Cann *et al.* 2012: 367).

Thus, the problem is not with the technique of minimalist idealisation but rather with the starting point. In order to account for the complete desired property or phenomenon, for instance anaphora in English, we need to consider a different minimal model, such as the discourse level or dialogue data. The modelling strategy remains constant in this case. DS modelling merely starts its idealisations from a different place but aims to generate the property of interest (grammaticality or anaphoric bind-

ing) in a minimal way. I think that *isolation* might be somewhat harder to establish as most DS models are quite integrative.

In DS, linguists are still in the business of developing minimal models and attempting to “generate” (in the sense of minimal generation) the properties or explanations thereof by offering a causal basis which ignores extraneous material. The framework might seem to aim for “completeness” but this completeness should not be confused for completeness at the initial modelling stage. In other words, the models aim to account for phenomena such as anaphora, ellipsis, quantifier scope etc. in a more complete way than their predecessors (of the generative persuasion) but they aim to do so through the most economical means possible (replete with the gamut of *ceteris paribus* modifiers and the like). For instance, underspecification plays an important role in the theory. This is a technique used to represent or generate multiple semantic (or other) representations within a single representation. Underspecification is a common technique for dealing with a wide range of ambiguities (both lexical and structural) in natural language semantics (and processing), without necessarily altering anything at the syntactic level.³ Semantic underspecification is basically an intentional omission of linguistic information from semantic description. The underlying idea is to postpone semantic analysis until it can be executed in such a way that various ambiguities can be resolved. In other words,

The key idea of underspecification is to devise a formalism which allows to represent all logical readings of a sentence in a single compact structure. Such a formalism allows one to preserve compositionality without artfully casting pure semantic ambiguities into syntactic ones. (Lesmo and Robaldo, 2006: 550).

This process amounts to a type of storage of interpretations without immediately checking for consistency. At a later stage these interpretations are pulled out or extracted and interpreted in parallel. A given semantic representation can be underspecified in one of two ways.

1. atomic subexpressions (constants and variables) may be ambiguous, i.e. do not have a single value specified as their denotation, but a range of possible values;
2. the way in which subexpressions are combined by means of constructions may not be fully specified (Bunt, 2007: 60).

These paths specify constraints on representations of meaning and are often viewed as meta-representations which display all the representations that satisfy the set of constraints, i.e. all the possible readings of an expression.

The point is that underspecification is a means of capturing multiple meanings within a single structure and a clear example of a minimalist idealisation. Underspecification is essentially misrepresentation for the sake of disambiguation at a

³Representing the scope ambiguities through alternative syntactic derivations can lead to an explosion of ambiguity and the need for innumerable alternative syntactic configurations. See Bunt and Muskens (1999) for a proof of this based on the ambiguity in an average Dutch sentence.

later stage in the process. The rules of the grammar (or DS) then apply to this idealised compact structure. Many of the other techniques utilised in DS are similarly motivated. The framework (non-trivially) exploits the strategy of minimal generation, as I have described in the previous chapter. In this way, it is within the modelling paradigm of the broader generative tradition in linguistics despite theoretical differences.

5.3 Model-theoretic Syntax and Overgeneralisation

Before concluding, I think it expedient to address a potential objection. One potential concern when offering accounts of modelling in linguistics (and the sciences in general) is the overgeneralisation of explanation. In describing a phenomenon which admits to certain vague or imprecise notions such as ‘causality’, ‘minimal’, ‘generation’ etc., a theorist can often provide explanations which trivially capture too much (or everything) and thus fail to distinguish between relevant alternatives. If all of linguistics from Hockett’s finite grammar to Smolensky’s harmonic grammar or Croft’s radical construction grammar could be explained in terms of minimalist idealisation, it would be no surprise that dynamic syntax followed suit. Fortunately, I believe that this is far from the case. In what follows, I will briefly mention some linguistic frameworks (or families of frameworks) which I believe do not have minimalist idealisation at their core. Once again, it is in no way my claim that minimalist idealisation is a preferable modelling strategy. I take no position on the fecundity of one type of modelling over another, my project is merely a descriptive one. I should also note that this analysis is not meant to be exhaustive. Many frameworks, including generative grammar, can and do involve other forms of modelling and idealisation. My claim is that minimalist idealisation is at the centre of many of these frameworks, not that it is the only strategy used to model linguistic reality within them (or to theorise more generally).

As I have mentioned, there are some similarities between both the parallel architecture and DS to model-theoretic approaches to grammar, I think that this would be a good point at which to describe these non-minimalist idealisation approaches. In Pullum and Scholz (2001), the notions of generative-enumerative versus model-theoretic syntactic formalisms are discussed and teased apart. The former are related to the formalisms discussed in the previous sections (with the exception of DS which has elements of both). These formalisms drew inspiration from the syntactic (or proof-theoretic) side of mathematical logic (and Post’s work on the subject)⁴. However, model-theoretic approaches were developed from the semantic side of logic and diverge from the generative-enumerative approach significantly. In this way, I think that model-theoretic syntax idealises its models in a distinct way as well, i.e. not *via* minimalist idealisation.

⁴See Part I.1.

One of the core technical notions of the previous formalisms was that of “generation” in the ‘recursively enumerate’ sense of defining a device with a finite set of rules capable of generating an infinite set of sentences/strings/structures. On the contrary “[a]n MTS [model-theoretic syntax] grammar does NOT recursively define a set of expressions; it merely states necessary conditions on the syntactic structures of individual expressions” (Pullum and Scholz, 2001: 19). Think of this approach in terms of model-theory. A sentence is well-formed iff it is a model of the grammar (defined in terms of constraints which act as the axioms of the formalism). To be a model of the grammar is to be an expression which satisfies the grammar (meets the constraints). Consider the first-order analogy. To be a model of arithmetic is to satisfy (or make true) the axioms of arithmetic (Peano or others). There are non-standard models of course and Gödel’s famous incompleteness result showed that there can never be a complete axiomatisation of arithmetic (i.e. no system will be able to capture all the truths of arithmetic). The point is that the idea of “being a model” of a grammar in this sense is quite divorced from the idea of “being generated” by a given grammar.

Formalisms such as Generalised Phrase Structure Grammar (GPSG) and Head-driven Phrase Structure Grammar (HPSG) are examples of this constraint-based approach.⁵ Some differences between the approaches involve concepts such as set cardinality. In section 4.3. we saw the motivation behind early generative grammar was to capture the notion of linguistic creativity described in terms of discretely infinite output. This is a corollary of the generative-enumerative approach, i.e. generative grammars generate or produce a fixed number or set of expressions (the upperbound is \aleph_0 or the cardinality of the natural numbers). Contrary to this, model-theoretic or constraint-based grammars do not impose such size-limits and are generally non-committal in terms of cardinality (which is not to say that they cannot account for creativity).

Now, there is certainly a parallel between specifying a set of axioms or constraints in defining a grammar and specifying a set of rules in generating one but the latter approach is more in line with minimalist idealisation than the former for important reasons. There are other significant differences but I shall focus on those that involve (or result in) a shift in modelling practices. We have already seen that the concept of infinity generation is abandoned on the model-theoretic approach but another important (relevant) departure from minimalist idealisation is the level at which the models are defined. For frameworks within this paradigm, models are individual expressions not sets of such expressions. In generative approaches, for the sake of generality, there was a movement towards categories (as sets) of ex-

⁵Although in the introduction of their textbook, Sag, Wasow and Bender (2003) speak in generative terms about the purpose and nature of grammar. “Thus we will again and again be engaged in the exercise of formulating a grammar that generates a certain set of word strings – the sentences predicted to be grammatical according to that grammar” (2011: 21). They also stress the importance of capturing the infinite set of expressions of natural language (through Kleene star operations and the like).

pressions and rules involving these categories (recall the X-bar rule schemata in section 4.3.1). Model-theoretic accounts quantify over specific expressions and the structures relating to these expressions.

For example, if trees are the intended models, quantifiers in the statements of MTS grammar range over a set of nodes, not over a set of trees (Pullum and Scholz, 2001: 23).

There are some important consequences (or perhaps advantages) of this feature of model-theoretic approaches. They all seem to be related to the greater specificity or accuracy which they allow the grammar to express or capture. For instance, expression fragments can more readily be treated under this framework. These could be fragments with syntactic structure or information (and semantic or phonological as well) that are not strictly grammatical (such as *and of the* example in Pullum and Scholz (2001)) and thus would not be generated by a generative grammar. By focusing on individual expressions we can also capture the use and proliferation of neologisms and the lexically creative aspect of natural languages. Given that the lexica of various languages are constantly changing, a formalism (or family of formalisms) which can capture (or at least not make bad predictions) about such expressions would be useful.⁶

There are a number of other such differences related to models as individual expressions. Some linguists (within the probabilistic school) have been claimed that the generative-enumerative approach can be described as “prescriptive” in the pejorative “grammar school” sense from which linguists have taken pains to separate themselves. There are constructions and phrases that pop up all over human language (and corpora) that would be deemed simply ungrammatical in the generative sense (i.e. not generated by any rule).⁷ Manning (2003) claims that the generative approach (what he calls “categorical linguistic theories”) are prescriptive in the sense that they place hard boundaries on grammaticality when these boundaries are much fuzzier in reality. This, however, is in keeping with idealisation. Nevertheless, the probabilistic linguistic models he suggest in their stead are beyond the current scope. Importantly, the criticism is related to the model-theoretic approach and its method of idealisation. In one way, generative grammars idealise too little (as *per* minimalist idealisation) and, on the other, they overgenerate (due to the lack of specificity). Hence the need for theta-grids and other ways of narrowing down the grammatical output of the grammar in generative frameworks such as Government and Binding.

⁶Both language fragments and neologisms can be captured by the incremental word-by-word parsing formalism of DS in a generative-enumerative way.

⁷This is related to the motivation behind DS and its claim that generative grammars offer incomplete descriptions since this miss out on relevant and systematic data found in corpora and multi-person discourse. I thank Geoff Poole for pointing out to me that a traditional GB/Minimalist account in terms of ‘ellipsis’ might work as well for these cases.

Constraint-based grammars focus on individual expressions and their satisfaction of a certain set of constraints. In so doing, they not only move away from the idea of minimal transformational derivation, but also admit for increased specificity at the individual expression level and in turn at the initial model level. The model aims to admit and satisfy as many constraints as needed to encode syntactic, semantic and other information, in order to characterise grammatical well-formedness. However, the descriptions involved in the models of the grammars can be quite complex. In fact, the models or expressions can be infinite in length on the constraint-based view (a welcomed feature for adherents of Langendoen and Postal's vastness proof). Whereas in generative grammar this is not the case (without significant modification on the notion of "generative" or "derivation"). This is another advantage of the MTS framework according to Pullum and Scholz (2001). So our base case could begin with a potentially infinite model with a multitude of constraints (well-motivated of course). With the development of generative approaches in sections 4.3 and 4.3.2 we saw a progression toward generality and the exclusion of (even causally relevant) material, in the model-theoretic framework of constraint-based approaches we see the reverse and a progression toward specificity and the inclusion of more information some causally relevant material, others not necessarily so (e.g. there might be irrelevant phonological information for instance). In addition, these approaches tend not to be modular or make use of *isolation* idealisation. For instance, feature structures, which can be thought of as functions from sets of features to values (or valences), in HPSG, are used to model grammatical categories as information structures. These structures can be extremely complex (as anyone who has used or seen a tree in HPSG can attest). The inclusion of semantics introduces additional information and structure into the features structures.

The richer feature structures we are now using, together with our highly schematized rules, have required us to refine our notion of how a grammar is related to the fully determinate phrase structure trees of the language (Sag *et al*, 2003: 167).

To close off this sketch, let us consider a basic operation in the grammar, namely unification. This operation takes two feature structures and creates one that contains all the information (and constraints) of both (as long as it is not inconsistent). Thus we are building larger and larger information structures into the scientific models of the grammar. I think we have moved quite a distance from minimalist idealisation and the modelling practices that come with it. One aspect of minimalist idealisation is that de-idealisation is very often not possible, recall the Ising model of ferromagnetism mentioned in section 4.3 in which particles are represented simply as points along a line. Feature structures and model-theoretic syntax allows for gradability of representation and the re-introduction of removed material. To be more specific, in formal language theory, sentences are modelled as semantically vacuous strings. A grammar, as a generative device, specifies the types of rules applicable to these

strings in order to generate different sets of strings, i.e. languages, of varying complexity. For example adding recursive rules to the rules which generate regular grammars gives rise to context-free grammars and so on (this picture is overly simplistic for the sake of illustration). Given this idealisation, there is simply no room for semantic content or phonological character to enter into the resulting model. These aspects of language are dealt with separately (in line with isolation idealisation). Contrasted with this, in constraint-based or model-theoretic approaches we can admit as much information into the base syntactic feature structures as we like, including semantic and phonological features. The models are not incompatible with introducing more and more information or features in order to “come closer” to the real world target system. This is the hallmark of *Galilean idealisation* (or mere abstraction) as Weisberg (2007, 2013) describes it or distortion for the sake of computational tractability with the possibility of reintroduction of abstracted or idealised material. However, this is generally not a feature of minimalist idealisation.

5.4 Subconclusion

It might be useful at this juncture to consider why these chapters should be of particular importance or interest within and outwith the field of linguistics. Linguistics is a scientific chimera, often lacking a clear unified methodology, theoretical persuasion or direction. The dominance of the generative tradition is receiving increased scrutiny and there is a plenitude of frameworks waiting in the wings to take its place. On the one extreme, divergences are often exaggerated and these frameworks are considered to be incommensurable (in the Kuhnian sense). On the other extreme, genuine differences are overlooked and considered to be mere ‘notational variants’ of one another (in the Chomskyan sense). The present chapters aimed at finding a middle ground in the identification of commonalities, in terms of scientific modelling practices, while respecting genuine theoretical advancements and divergences. In the conclusion of this part, we will have something more to say about these issues.

Here, however, I have argued that the generative tradition encompasses two related varieties of modelling, namely minimal generation and isolation. These modelling strategies fall under the auspices of a modelling practice commonly found in the sciences, namely minimalist idealisation as described by Weisberg (2007, 2013). Under this paradigm, linguists aim to identify a core (minimal) causal model which gives rise to a property or phenomenon of interest while ignoring other (even potentially relevant) features of the target system. In this sense, the strategy involves a distortion or idealisation of the target system in order to capture the least set of elements responsible for a given property *via* the initial misrepresentation. Following a line set by Blutner (2011), I extended this analysis beyond the standard accounts within generative grammar such as Government and Binding and the Minimalist program, to include Jackendoff’s parallel architecture and optimality theory.

Lastly, I attempted to unite the modelling practices of the generative tradition with a competing approach which lacks the similar theoretical underpinnings of the parallel architecture and OT, namely the dynamic syntax of Kempson *et al* (2001). I argued that although the theoretical underpinnings of this latter framework do seem distinct (and genuinely are) from those of the larger generative programme, they approach the target system of natural language in similar ways. For the sake of nontriviality, I presented an overview of linguistic frameworks which do not share this modelling approach.

Chapter 6

Modelling in Semantics

Formal semantics is a mathematical framework for the investigation of meaning in natural language, its structure and significance. As a field, semantics has been beset with objections ranging from claims as to its triviality, its limited scope and its questionable scientific target. In this chapter, I present a view of formal semantics as applied mathematics where the “application” maps formal syntactic models onto the world or assumed structure of the world. I attempt to address the questions of the purpose, target and efficacy of formal semantics through the analogy with applied mathematics and also thereby shed light on its explanatory and predictive power and the syntax-semantics interface. The account presented here draws from detailed comments on the nature of compositionality, negative polarity items and underspecification techniques in semantics.

6.1 Formal Semantics as Applied Mathematics

In the previous chapters, I aimed to provide an account of the modelling practices inherent in specific theories of syntax and redefine the generative programme in modelling terms (or in accordance with the first grade of involvement of Part I). In this chapter, I aim to pair that analysis up with the modelling involved in formal semantics. It will not, however, be my mission to analyse individual theory development in this chapter as it was in the last. Firstly, because the modelling perspective seems to have been more generally acknowledged by semanticists of various persuasions and secondly because my purpose here will be to offer an underlying philosophical characterisation of modelling in semantics in accordance with a tradition within applied mathematics and its philosophy.

A good place to begin to understand the nature and place of formal semantics in linguistics and philosophy is by appreciating its essential connection to a theory of formal syntax. Although the mathematics which inspired formal syntax (the generative enumerative formalisms of proof theory) diverges from the model theoretic

underpinnings of formal semantics, the scientific import of semantics can and should remain substantively linked to syntactic theory and phenomena. In this chapter, I argue that formal semantics can be construed as the mathematical application of model theoretic techniques to a formalised structure of grammatical discourse given to us by theoretical syntax. In this sense, formal semantics is an applied mathematical discipline whose target is not necessarily the empirical world directly but rather the structures generated by a syntactic analysis of natural language, i.e. syntactic models.

An alternative approach to understanding the nature of semantics can be gleaned from the tradition connected with a theory of truth. Davidson (1965) famously reversed Tarski's direction of explanation for truth as relying on a fixed theory of meaning and instead used a fixed theory of truth to establish a systematic theory of meaning.¹ Unlike syntactic models which aim to capture facts about grammaticality capable of construal independent of correspondence to objects or situations in the world, a largely structural enterprise, semantic theories, on this view, tend to attempt a link between linguistic formalism and nonlinguistic reality. Thus, it is no surprise that semantics has held the attention and fascination of philosophers of language to a more significant extent than its syntactic counterpart.² There are many insights to be had within this framework. However, we have to be wary of conflating the scientific project of modelling linguistic phenomena with foundational projects aimed at accounting for them philosophically.³

The rich philosophical tradition on linguistic meaning notwithstanding, formal semantics has emerged as a programme strongly connected to the syntactic project of linguists. I will continue this approach to understanding the role and nature of semantics in the present work. I will further argue that in order to appreciate the scientific explanations present in semantics, i.e. how semantic modelling works, this view is essential.

Semantics as a distinct scientific discipline within theoretical linguistics has had a controversial history and continues to be divisive among some theorists. Many linguists, following Chomsky, believe that a theory of meaning has no part in an explanation of linguistic competence or the faculty of language narrowly construed (see Chomsky, Hauser and Fitch 2002). Others, such as Horwich, advocate a more

¹"It is possible to view a Tarski truth characterization for a language L as simply specifying the extension of 'true' for L , explaining how the truth value of a sentence depends on the semantic properties of its parts, and providing the basis for accounts of logical truth and logical consequence" (Soames, 1984: 416).

²This aspect of semantics is related to the notion of ontological commitment. Rayo (2007) makes the connection between truth conditions and ontology by stating that "[t]o describe a sentence's ontological commitments is to describe some of the demands that the sentence's truth imposes on the world" (428). These demands are captured by the truth-conditions.

³Yet another distinction is presented by the recent introduction of "meta-semantics" or the philosophical investigation of the purpose, place and methodology of semantics within scientific inquiry. I would consider the present chapter a part of this nascent tradition.

Wittgenstinian line in terms of usage-based concepts. A growing trend in the philosophy of language, which includes radical contextualism, aims to narrow the range of phenomena which is in need of specific semantic treatment (as opposed to say, treatment in terms of pragmatics). Some, such as Sampson (2001), claim that linguistics dances on the boundary line between the sciences and the humanities, whereas syntax can receive a scientific (read: empirical) treatment, semantics is significantly more amorphous and falls squarely on the humane linguistics agenda.

In section 6.2, I investigate the aim, target domain and methodology of semantics. In section 6.3, I argue that semantic analyses act as a sort of engineering which makes our syntactic models applicable to the world and linguistic communication by interpreting them. I will argue that formal semantics should be construed as a type of applied mathematics in the spirit of Hughes (1997), Suárez (2004) and others, much like engineering or mathematical physics, in which the formal model theoretic apparatus is applied to the syntactic models of human language use and cognition as interpreting structures or mappings. First, however, an outline of the aims, target and methods of formal semantics is in order. In section 6.4, I use the case of NPI licensing to investigate the predictive capacity of semantic models. Lastly, in section 6.5, I use the philosophical apparatus set up in the previous two sections to account for the syntax-semantics interface.

6.2 Aims, Target and Methods

6.2.1 Semantic Explananda

The literature on formal semantics commonly presupposes a modelling perspective. Semanticists often refer to their task as modelling semantic phenomena or creating semantic models in order to capture semantic facts. For instance, Yalcin claims that “[l]inguistic syntax and semantics are ultimately concerned to produce explanatory models of aspects of the knowledge that underwrites this capacity” (2014: 36). In a discussion of context-sensitivity, Dever (2012) asserts “[c]ontext-sensitivity of language can be modelled in a formal semantic theory by assigning semantic values relative to a context, or by assigning semantic values that are functions from contexts to more standard semantic values” (55). Here the optional nature of the modelling device indicates that the semantic model is an indirect representation of the target property (e.g. context-sensitivity). Though modelling parlance is widespread, this by itself is not enough to establish that semantic theory constitutes scientific modelling. In fact, the term “model” itself is ambiguous between the concept of a certain interpreting structure in model theory (consisting of a domain of objects and interpretation function) and a scientific model (a special type of representing device). Needless to say, the influence of model theory on formal semantics did not aid in the resolution of this confusion. As Hodges puts it,

Semantic models don't belong in the same list as scale models, analogue models, diagram models, mathematical models and so forth. The reason why they are called 'models' at all is historical and rather indirect (2009: 667).

Thus, there is a distinction between semantic models in the mathematical logic sense and semantic modelling in the scientific sense, i.e. the modelling practised in formal semantics. We would do well to keep these different conceptions in mind going forward. Nevertheless, *pace* Hodges, modelling in semantics might, and I argue does, resemble modelling in the applied sciences. However, in order to understand modelling in any scientific domain, one needs to isolate the aim, target and method employed to achieve the modeller's aim.

On to the first task. Semantics is often considered to be in the business of articulating a theory of meaning. The theory of meaning in question has certain important characteristics. For Davidson, one core explanandum was learnability which he took to be evidence of compositionality. "Then we may state the condition under discussion by saying: a learnable language has a finite number of semantical primitives" (Davidson, 1965: 9).

This coupled with the ubiquitous claim that there are an infinite number of natural language sentences results in a particular version of the principle of compositionality.

When we can regard the meaning of each sentence as a function of a finite number of features of the sentence, we have an insight not only into what there is to be learned; we also understand how an infinite aptitude can be encompassed by finite accomplishments (Davidson, 1965: 8).

Although learnability is certainly a worthy aspect of linguistic investigation, it does not often fall under the purview of formal semantics explicitly (although it is assumed).⁴ Rather, the aim of semantics is to account for "semantic facts" such as they are "found in the world" so to speak (represented in (6.1) to (6.5) below). This aim is not to be confused with the explanation of meaningful linguistic behaviour in general which would include pragmatic phenomena.

Stating the linguistic facts for which semantics is responsible is not an easy task. The problem is that even selecting what counts as data is a theoretical matter, especially in linguistics. For instance, generative linguists often strive for what they call explanatorily adequate grammars which take language acquisition data as explananda. In Dynamic Syntax, grammars are influenced by parsing phenomena (see Part II.5). Therefore, deciding which linguistic phenomena make it into semantic models is not a theory neutral activity. The best one can hope for is to capture

⁴This division of labour corresponds to similar adequacy conditions for a theory of grammar described in Chomsky (1965), where the level of *explanatorily adequate grammars* aim to account for language acquisition data.

the intersection of various facts which often receive treatment from formal semantic theories.⁵ A good place to start is with a canonical textbook of Larson and Segal.⁶

Larson and Segal (1995) split the data in need of formal semantic description into five categories, primary or “immediate” facts, secondary ambiguity facts, semantically anomalous statements, entailment and thematic relations respectively.

(6.1) Snow is white.

(6.2) Flying planes can be dangerous.

(6.3) *The theory of universal grammar stole my sandwich.

(6.4) 1. Susan is a woman.

2. Susan is a mammal.

(6.5) 1. The train travelled from London to Paris.

2. The inheritance passed from John to Mary.

3. The substance changed from liquid to gas.

The first three of these examples represent a class of expressions which are commonly associated with formal semantic description, namely standard declaratives involving predication (and their disquotational counterparts), ambiguous sentences or sentences with two or more readings and semantically anomalous sentences (with otherwise acceptable syntactic structure).⁷ The next set of sentences in (6.4) are connected by some kind of entailment relation. And finally the last set of sentences are thematically related in that they all convey the action or concept of change or transition. Thematic relations are usually linked to the notion of argument structure. In Jackendoff (2002), thematic role determines syntactic word order *via* a principle of linking hierarchy for NP arguments. He claims that the principle belongs neither to syntax nor to semantics proper but rather the interface between them.

⁵We can see the difficulty in demarcating the “neutral” data if we consider examples of the so-called “non-substitutivity of logical equivalents” from Partee (1979). Formal semantics treats the pair *Irene believes that p* and *Irene believes that q* differently even when *p* is equivalent to *q*. This suggests that human psychology plays a role in semantic modelling.

⁶This approach is not merely a convenience. I follow Giere (1988) here in starting from an analysis of textbooks as opposed to active research in a field (I do not maintain his focus on them throughout though).

If we wish to learn what a theory is from the standpoint of scientists who use that theory, one way to proceed is by examining the textbooks from which they learned most of what they know about that theory (63).

⁷The latter phenomenon has recently been claimed to have significance for formal semantics and lexical semantics in general with relation to the type theoretic structure required within the formalism. See Asher (2011).

In all of the above cases, linguistic meaning is associated with a particular form or syntactic structure whether it is subject-predicate or VP, quantifier-pronoun or thematic role. Grammaticality or the absence of grammaticality as in the case of sentences like (6.3) provides a guide to semantic interpretation. This is generally not the case with pragmatic phenomena. Consider the following examples said of a job applicant in a reference letter and to a friend by a second language English speaker respectively.

(6.6) Tolani has excellent handwriting.

(6.7) *She said me to stay.

In the first of these cases, more than syntactic arrangement is needed in order to interpret or retrieve the intention of the speaker, such as context, tone etc. In (6.6) the sentence is grammatical yet the speaker or writer does not intend the literal meaning of the sentence. (6.7) is not a sentence of English, yet it is intelligible. Formal semantics is not responsible for these sorts of phenomena.⁸

This layout of semantically relevant data dovetails in some respects with Yalcin's (2014) claims as to the nature of what constitutes a semantic datum. His characterisation includes productivity facts, entailment facts, communication facts, acceptability facts and truth or appropriateness facts. Productivity facts are supposed to be related to the ability to produce and understand previously unencountered linguistic expressions (what I call "compositional semantic facts" below). Entailment facts are related to implication, consistency and contradiction (e.g. sentences in (6.4)). Communication facts are related to the systematic transfer of information within any language user's arsenal. Acceptability facts are related to anomaly (as in (6.3)). And finally truth or appropriateness facts are related to the fact that we judge sentences like (6.1) as true or appropriate in certain contexts.

Most of these types of data are friendly amendments to the view currently under discussion. However, I would hesitate to include "communication facts" under the remit of formal semantics proper for a few reasons. For one thing, communication facts are too broad and encompass all kinds of pragmatic and even nonlinguistic information transfer. Furthermore, "systematic information transfer" need not involve grammatical form of any sort. Incompetent speakers of a language can still convey information in a systematic manner, including this data might not serve us well as semanticists. Consider a pidgin language established for trade purposes between two remote villages. Although standard pidgin languages involve significant syntactic poverty, they still manage to convey systemic information which includes plural and number specifications (often through reduplication as opposed to derivational

⁸In fact, a criticism of the generative semantics movement of the early 70's, one which is associated with its downfall, is that it attempted to model too much phenomena. See Newmeyer (1996) and Ludlow (2011).

morphology). There are no native speakers of pidgins and thus the semantics of such languages resemble a hodgepodge of resources and contextual inferences.⁹

Therefore, the way in which to weed out extraneous factors, such as pragmatic phenomena or ungrammatical but otherwise meaningful signs, is by returning to a point suggested by Davidson, mentioned earlier. The difference between data such as that represented in (6.1)–(6.5) and (6.6)–(6.7) (and some communication data) above is that the meaning of the former can be derived compositionally from the syntax while the meaning of the latter cannot.

Thus, the principle of compositionality is not only a means of capturing desirable linguistic properties such as learnability, productivity and systematicity but it also acts as a filter for relevant versus irrelevant data for linguistic semantics. More on the nature of compositionality as a modelling choice in section 6.5. Higginbotham describes the overarching point in characteristically illuminating fashion, “[s]emantics connects forms with meaning: but the theory of meaning and the theory of form do not proceed along separate tracks. Information about each can provide evidence for the other” (1989: 159).

The picture that is emerging of the “semantic facts” for which a scientific theory of meaning should account is based on the relationship between syntax and semantics. On this view, formal semantics aims to explain only “compositional semantic facts”. Or in other words, semantic analyses do grow on trees (or directed graphs or categorical proofs etc.).

6.2.2 Target of the Model

Even if the so-called “semantic facts” or domain of inquiry are agreed upon pretheoretically, there are divergent views on what exactly the overarching target of semantic modelling is, i.e. what explaining or modelling the aforementioned semantic facts amounts to. Another way of distinguishing aims and targets is that the aim could be the explanation of compositional semantic facts but the target could be considered in terms of the location of those facts.

Cartesian Semantics

On the dominant tradition, linguistics is the study of a distinct language faculty inside of the mind/brain of the language user. Following Higginbotham (1989), this view can be referred to as “Cartesian linguistics” and correspondingly “Cartesian Semantics”. On this view, the target of linguistic inquiry is a type of tacit knowledge on the part of speakers brought out by their linguistic judgements. In an almost

⁹These languages are learnable but cannot be acquired. When pidgins do get adopted by linguistic communities, they tend to undergo what is referred to as “creolisation” which involves the establishment of systematic syntactic and semantic rules. Nevertheless, this latter point is not uncontroversial.

Kantian move, the collective idiolects of speakers constitute an objective reality and subsequent target of linguistic inquiry which in this case is an internally represented system of rules or I-language. The idea, initially proposed by Chomsky (1957), is that a grammar (or generative grammar) of a language L constitutes a scientific theory of an internalised rule system for generating all and only the grammatical sentences of L . A speaker ‘knows’ or ‘cognises’ this system of rules and produces judgements or linguistic intuitions based on this knowledge of the grammar. However, the extension of this idea to semantics proved controversial for most Chomskyans (and in part resulted in the Semantic Wars of the 70’s). Nevertheless, many contemporary semanticists consider themselves to be squarely rooted in the generative tradition.

For example, compare the following quotes from Yalcin and Glanzberg respectively on the object of semantic theory.

I shall also assume, as is generally assumed in the tradition of generative linguistics, that linguistic competence is the primary subject of study for linguistic theory. Thus, in effect, linguistic theory studies what knowledge—or more generally what cognitive states—underwrite our linguistic abilities (2014: 261).

and,

From the perspective on semantic theory I am recommending here, communicative uses of language reveal aspects of the state of mind which consists in understanding and speaking the language. It is the state of mind, or the semantic aspect thereof (‘knowledge of meaning’, understood in the technical sense), that semantics is foremost concerned with modeling (2014: 31).

Pinning down precisely what sort of “knowledge” is involved in the exercise of linguistic abilities is a notoriously difficult task (but as Glanzberg notes, does involve “modeling”). Following Devitt and Sterelny (1989) and Devitt (2006), let us take G to be a grammar or set of rules the members of which form a fragment of English. The question is then, what relation does a speaker of English have to G ?

There are few options. We could say that a speaker behaves in accordance with G (“as if governed by G ” in Devitt and Sterelny’s words). However, this will not do for the Cartesian linguist since the speaker could be behaving in accordance with some internalised G' which is identical to G in behavioural output (see Part I.3.1). The next possibility is that the speaker ‘knows’ that G is a fragment of English grammar and applies it accordingly. This is the strong representational thesis (of Devitt 2006). It attributes propositional knowledge (knowledge-that) of the rules of a language to the language-user. Now of course, the speaker does not possess explicit propositional knowledge of her language, but rather this knowledge is tacit.¹⁰ Unfortunately, this does not explain much. In attempting to explain what exactly this

¹⁰The distinctions between explicit, implicit and tacit are not always stated clearly. A rough characterisation (along the lines of Dummett) is that explicit knowledge is verbalisable upon re-

tacit knowledge constitutes, we are returned to the same issue as above, i.e. does it involve behaving in accordance with G or having a(n) (unconscious) propositional attitude towards G and how do we distinguish its rivals G' etc.?

The challenges for this view do not cease here. Chomsky (1980a) introduced the term of art “cognise” in many ways to avoid the sorts of concerns that are brought forward with the term “knowledge”. It is unclear how this strategy eliminates the worries though.¹¹

Public Language

A rival conception of the target of semantic (and linguistic) modelling is usually presented under the umbrella term of “public language”. A chief proponent of this view is Michael Dummett. For Dummett, natural languages are constituted by the social practices in which they are used. Thus, the target of semantics is the public languages which emerge from shared linguistic activities in a community of language users. These shared practices are explained in terms of conventions.¹² Millikan (2005) similarly argues that “a primary function of the human language faculty is to support linguistic conventions, and that these have an essentially communicative function” (25). She thus offers an alternative approach to the Chomskyan view that there is no “public object” which corresponds to a natural language. However, contrary to many views on the nature of a public language, such as the Platonism of Katz and Postal, Millikan argues that,

[t]he web of conventions that forms the mass that is public language is not an abstract object but a concrete set of speaker–hearer interactions forming lineages roughly in the biological sense. These lineages and their interactions with one another are worthy of scientific study. Nor are their properties derivative merely from the properties of I-languages (2005: 28).

Non-cartesian views on the target of linguistic theory emphasise the communicative social aspects of natural language. A specifically semantic proposal in this vein is contemporary Brandomian inferentialism. The core idea of this framework is that in the same way that the meanings or content of logical constants can be determined by the inferential roles they play in a logical system (*via* introduction and elimination rules), the meanings of ordinary terms and words in natural language

quest/prompting, implicit knowledge is derivable from explicit (through inference or something of that sort) and tacit knowledge is something like behaviourally attributable but not verbalisable or derivable (this is vague). On other accounts, tacit knowledge has been described as I have described implicit knowledge here.

¹¹See Ludlow (2011:50) for a discussion of why knowledge is too “thick” a notion for our linguistic purposes and cognise is too “thin” a notion.

¹²See Davidson (1986) for the claim that knowledge of shared linguistics conventions is neither necessary nor sufficient for communication. Also see Lepore and Ludwig (2007) for overview of the debate between Davidson and Dummett on this issue.

are determined by their inferential roles. This is a radical idea which challenges a number of tenets of the referential/truth-conditional theories of meaning the likes of which have dominated the semantic landscape since the work of Frege and Russell in the early 20th century.

In a recent exposition of the inferentialist programme, Peregrin (2015) argues for the central place of rules within semantics. Unlike the Chomskyan conception of rules which are internal and individualistic (two aspects of the “I” in I-language) and the external communicative aspects of natural language are purely ancillary in most respects, rules are norms established by the practices of linguistic communities.

To be sure, if an expression has a meaning within a linguistic community, then the speakers of the community *will* conceive of it in certain specific ways. However, this is not enough to establish the fact that it means what it does. An essentially private act of conception is not capable of grounding the essentially public institution of language. That people of some community mentally associate the word *spider* with a certain kind of animal is a fact of their individual psychologies not capable of establishing the fact that *spider* expresses, within their language, the concept of *spider*; what is needed alongside any private associations are some public practices that make the link between the word and a concept public and shared (Peregrin, 2015: 44).

Although the concept of rule remains central on this view, the order of explanation is reversed from the inside-out to a view of meaning being determined from the outside-in, from external social normative practices to internal mental representations. The target of formal semantics is then more sociological than psychological.

These debates are certainly of great interest to the study of meaning in general and the theories that are informed by such a study. However, I do not think that they necessarily affect formal semantics from a modelling perspective. Models are indirect representations of a target system, in this case natural language. In this way, models are compatible with a number of divergent views on the ontological nature of the target and philosophical disposition of theorists. This is not to say that modellers are in no way influenced by the intellectual crazes of the day but the models are to a large extent independent of theories. As we have seen, semanticists of contrasting views tend to agree on the semantic facts (or “compositional semantic facts”) in need of explanation. Whether they think these facts stem from knowledge of internal rule systems of a language faculty or from external social practices of linguistic communities, the facts remain the same. The underlying theoretical target or persuasion of the modellers in this respect is largely orthogonal to the task of explaining, representing and predicting compositional semantic facts through the application of formal methods.

Autonomy of Models

There is a certain element of instrumentalism involved in the process of scientifically modelling a specific phenomenon. The autonomy or partial autonomy of models is well-established in the literature on the subject.

It is common to think that models can be derived entirely from theory or from data. However, if we look closely at the ways models are constructed we can begin to see the sources of their independence. It is because they are neither one thing nor the other, neither just theory nor data, but typically involve some of both (and often additional 'outside' elements), that they can mediate between theory and the world (Morrison and Morgan, 1999: 11).

This is to say that models are often constructed or built up in a piecemeal and heterogeneous fashion as argued by Suárez and Cartwright (2008) with relation to the Londons' model of superconductivity in physics. In addition, models tend to function in autonomous ways as well. In the case at hand, the same instrument (semantic model) can be used for various functions and to represent various targets both psychological and social. For instance, Klein (1999) claims that historically certain chemical formulas functioned as models and affected theory change similarly to Fitzgerald's pulley and rubber model of aether used to correct Maxwell's electromagnetic theory. There are many more such examples from the history and practice of science.

The point is that models, and in this case semantic models, can be assessed independently of theoretical targets and constraints. The view of formal semantic modelling developed here is one which takes no position on whether semantic knowledge or social practice is the correct theory of meaning in natural language.

Models can be used to study a single target, a cluster of targets, a generalized target, or even targets known not to exist. One can even engage in the study of a model without any target at all (Weisberg, 2013: 74).

This is not to say that, in every case, the relationship between a target system and a model can be explained independently of the modeller's intentions. Hodge's (2009) seems to hold the contrary view, namely that a model M and a system (or target) S "do involve a correlation linking M and S . But there is no need for them to mention anybody intending this correlation" (667). I think that this is a mistake. Consider the case of mathematical models with identical mathematical structures such as a given pendulum and a particular circuit or a spring. Without some account of the "construal" or interpretation of the model, there would be no telling these models apart. This is the view held by Weisberg (2013), in which "we can say that these

models share a common core mathematical structure, and what differs are theorists' construals" (72).¹³

In Part I we discussed Lewis' (1975) alternative to both views described here. For Lewis, a language is both an abstract object, a function mapping forms to intensions (or meanings), and a social construct used by a linguistic community. The best way to describe the relationship between these two distinct aspects of natural language is that the language L used by a specific community c is a model determined by the conventions, defined in terms of regularities, of that community. On Lewis' picture, a language is comprised of both form or syntax and meaning or semantics. The language L is a function and within its domain are syntactic objects while its range is comprised of meanings. Once again, syntax and semantics form part of an ordered pair called a language. Higginbotham (1985) critiques this view of semantic facts being defined purely in terms of the range of linguistic functions. To this effect he has a modal argument stating that nonsentences or sentences not generated by the syntactic formalism (grammar) of a specific language should fall within the remit of semantics in addition to well-formed sentences. For Higginbotham, sentences and nonsentences are structurally ambiguous. I think that this emendation takes us too far a field from tractable compositional facts such as those presented in section 6.2.2.¹⁴

The common thread in all of these various frameworks is the focus on interpreted form. For Chomskyans, form or syntax is the central component of the language faculty and our linguistic competence. Higginbotham, Glanzberg and Yalcin¹⁵ extend this account of linguistic competence to include semantic knowledge. Semanticists of the inferentialist persuasion follow Dummett in the acceptance of a public language as a legitimate focus of scientific theory. However, they do maintain that sentences are the vehicles of semantic expression. In fact, Brandom stresses the "top-down" or sententialist aspect of semantic inferences and the reliance on compositionality, i.e. semantic analysis starts with whole sentences where the content of a sentence is determined by the role the sentence plays in a language game of "giving and asking for reasons" or it is "a matter of being able to play the role both of premise and of conclusion in inferences" (2007: 654).

Thus, the target system of formal semantic models are syntactic grammars or models. Whether linguistics is psychology, sociology or a hybrid theory does not

¹³A construal is comprised of an assignment, intended scope and two fidelity criteria for Weisberg, where the latter are "the standards that theorists use to evaluate a model's ability to represent real phenomena" (2013: 76).

¹⁴For one thing, this seems to be a similar move to the UG hypothesis of syntactic theory in which individual variation in form is assumed to be subsumed by a more universal patterning. UG is tendentious in syntax, in semantics it is almost untenable. To say that formal semantics needs to account for an underlying semantic structure invariant under different natural language instantiation is not only controversial but runs against the simplifying agenda of most modelling practices.

¹⁵Larson and Segal (1995) and Heim and Kratzer (1998) would also fall within this tradition.

directly concern the modelling practices of semanticists. The only way in which these theoretical positions would have mattered more to the current debate would have been if psychological or other social scientific methods were used in the construction of semantic models or formal grammars in general. In the next section, I aim to show that this is far from the case.

6.2.3 Methods and Model Theory

The methodology of formal semantics is well-documented through the means of various standard textbooks on the subject from the aforementioned Larson and Segal (1995) and the more generative linguistics inspired Heim and Kratzer (1998) to the more theory neutral GAMUT (1991). The common apparatus of formal semantic theories is familiar from formal logic, type theory and lambda calculus. These theories are usually built up from the denotation of individual morphemes or words, the terminal elements of syntactic trees, to complex elements conjoined semantically through functional application. As Dever (2012) puts it,

There is no precise delineation of what counts as formal semantics. Roughly, though, formal semantics is the attempt to give precise accounts of the relation between syntactic structures and semantic values, typically while making use of tools from mathematics and logic (49).

Two basic insights characterise the field, namely that the meaning of an expression is provided by the conditions under which it is true (at least for declarative sentences) and that meanings are composed according to the principle of compositionality, a version of which is presented below.

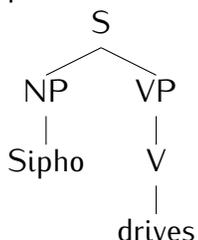
The meaning of a complex expression is a function of the meanings of its constituents and the syntactic rule used to combine these constituents.

There are a number of syntactic notions present in the above definition, such as “constituent” and “syntactic rule”, each of which is defined in a myriad of ways by various syntactic theories.¹⁶ The most common way of representing the relation between syntactic elements and semantic interpretation is by means of a homomorphism between two abstract algebras (one corresponding to the syntax and one to the semantics). The above principle can then be functionally (homomorphically) defined in the following way.

For every syntactic rule $\sigma \in \Sigma$ there is a corresponding meaning operation r_μ such that if $\sigma(u_1, \dots, u_n)$ is assigned a meaning, then $\mu(\sigma(u_1, \dots, u_n)) = r_\mu(\mu(u_1), \dots, \mu(u_n))$.

¹⁶In some formalisms, such as dependency grammar, there is no clear notion of constituency.

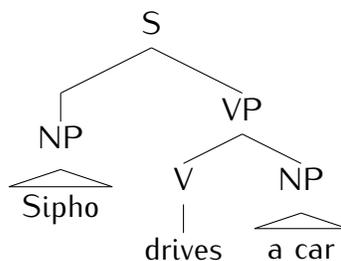
Formal details of the way in which this programme is exacted algebraically is beyond the scope of the present work. I will, however, present a sample formal semantic analysis in order to display the various mathematical techniques in the service of these models. For the purposes of exposition, I will simplify matters by sticking to an extensional truth-conditional semantics.¹⁷ Let us start with a basic intransitive sentence *Sipho drives* represented in the tree below.



Proper names and other physical objects are represented by the type e for entity and intransitive verbs are functions from entities to truth values $\langle e, t \rangle$ where the semantic value of a sentence is t . Following Heim and Kratzer (1995: 17), the semantic value of the nonterminal V or N node is inherited from the semantic value of its terminal counterpart. The interesting part of the analysis occurs at the non-terminal nodes which mark syntactic composition. In this case the S -node of the tree. By means of Fregean insight, the corresponding semantic rule is provided by functional application, where the semantic value of *Sipho* (indicated by the double line delimiters) is applied to the semantic value of *drives*.

- $\llbracket \textit{Sipho drives} \rrbracket = \llbracket \textit{drives} \rrbracket(\llbracket \textit{Sipho} \rrbracket)$

The entire sentence then has a semantic value t while the intransitive construction has a value represented by a function from objects to truth values or $\langle e, t \rangle$. Let us add a further level of complication by considering a transitive extension of the sample sentence. Consider the following tree.



The set theoretic object associated with this tree involves two levels of functional application (technically three if you include the determiner phrase).¹⁸ The semantic

¹⁷The mathematics involved in the intensional variant is very much the same as the extensional theory with the exception of the possible world type s and functions involving it.

¹⁸Most current accounts in syntax use DPs or determiner phrases instead of NPs even when there

value of a transitive verb has the type $\langle e, \langle e, t \rangle \rangle$ or functions from entities to functions from entities to truth-values represented below in two steps of functional application.

- $\llbracket \text{drives a car} \rrbracket = \llbracket \text{drives} \rrbracket(\llbracket \text{a car} \rrbracket)$
- $\llbracket \text{Sipho drives a car} \rrbracket = \llbracket \text{drives} \rrbracket(\llbracket \text{a car} \rrbracket)(\llbracket \text{Sipho} \rrbracket)$

The verb *drive* is ambiguous between a transitive and intransitive constructions (or optionally selects for direct objects). Modelling transitive verbs sometimes seems to require that we use more than just functional application, we might also use the technique commonly known as *currying* a function in mathematics (Heim and Kratzer prefer the term “Schönfinkelization”). Currying a function involves taking a function with multiple arguments or a domain of ordered n-tuples and turning it into a sequence of functions (or complex function), taking each argument individually. For instance, in the case of a ditransitive verb like *give*, instead of taking it to involve the Cartesian product of three entities to a truth-value (or $\langle e \times e \times e, t \rangle$), we take each in turn as in $\langle e, \langle e, \langle e, t \rangle \rangle$. Currying or Schönfinkelization, however, might not just be artifacts of the modelling. If we take semantic values as aspects of the target system and certain kinds of functions to represent these in our models, we could allow for the possibility that semantic values *qua* functions come pre-“curried” to a certain extent (I will return to the issue of direct modelling strategies in section 6.5.2).

Technically, transitive constructions can be handled within the bounds of functional application. What does pose a problem for strict functional application is phenomena such as object position quantifiers as in the examples below.

(6.8) Kgalema likes every politician.

(6.9) Nokuthula met some people.

The problem is that there is no straightforward way to apply functional application in this case since the semantic value of a transitive expression is type $\langle e, \langle e, t \rangle \rangle$ and the quantifier expression is of type $\langle \langle e, t \rangle, t \rangle$. One way in which to account for such data is by positing quantifier raising or movement rules. In essence, *every politician* or *some people* moves to a higher node in the tree and leaves a trace behind, thereby allowing the semantic values to be computed as before (*as if* the quantifiers were in the higher position). Of course, this analysis does require something akin to LF or logical form for surface forms to map onto. We will return to the issue of syntactic reinterpretation based on semantic considerations in the forthcoming sections. For now, however, the point is that unappended functional application is not enough to account for every syntactic combination.

is a null determiner place. I ignore these details for the sake of simplicity. Similarly for I(nflectional) P(hrase) and C(omplementiser) P(hrase) categories.

Lastly, before concluding the present sketch, we need to mention the use of λ -notation ubiquitous in formal semantics. We have already seen that in the process of mapping syntactic structures onto set-theoretic objects, we encountered complex functional notation and with the inclusion of quantifiers, reflexives and traces the notational complexity only grows. In order to limit confusion and name various set theoretic objects, we introduce λ -notation with variables from formal logic. Thus, a simple sentence like *Thembie works* becomes $\lambda x(x \text{ works})(\text{Thembie})$ where we replace the name *Thembie* with a variable x , add the name as an argument, and attach the λ -operator to bind the variable with the result that the λ -formula now represents a truth-conditional function. The formula for lambda abstraction used by semanticists is the following schema " $\lambda\alpha : \beta.\gamma$ " where α is a variable place reserved for the argument, β specifies the domain of the function and γ assigns a value for the argument α . The example in Heim and Kratzer (1995: 34) represents the successor function in arithmetic.

- $F_{+1} := [\lambda x : x \in \mathbb{N}. x + 1]$

The above definition can be read as the function which maps every x such that x is in \mathbb{N} to $x + 1$. This notation proves extremely useful in denoting various complex set theoretic objects corresponding to the semantic values which in turn represent syntactic objects and their composition. With the tools of functional application, type theory and λ -calculus, formal semantic modelling can represent the semantic values of a range of syntactic objects both simple and complex.

The essential aspect of the apparatus described so far is that it serves to further establish the scientific modelling nature of the enterprise of formal semantics. Jacobson (2014) explains the use of the aforementioned formal techniques in the following way.

[T]he language used to name model-theoretic objects is of no consequence for the theory and is only a tool for the linguist (the grammar itself simply pairs expressions with model-theoretic objects). But once the objects get sufficiently complex (e.g., functions with functions as their domains as well as co-domains) it is useful to have a simple way to write out these meanings without excruciating and/or ambiguous prose, and the lambda calculus allows us to do just this [...] a reminder that these name actual model-theoretic objects and that the formulas have no theoretical significance (Jacobson, 2014: 134).

Lambdas, types and variables are not found in nature.¹⁹ The semantic models we use to represent the meanings of syntactic structures operate in (partial) autonomy from what we take meaningful discourse or natural languages themselves to be, i.e.

¹⁹Pronouns might be thought of as the natural language equivalent of variables. For a contrary view, see Collins (in preparation) on variable involvement in syntax. Variable-free semantics is also a lively research programme.

states of the language faculty or patterns determined by social norms. In addition, the models abstract away from other linguistic data which might be connected to meaningful discourse or interpretation, such as phonological or pragmatic information. Furthermore, semantic models are a subset of more general mathematical models found in the other applied sciences such as mathematical physics and engineering. An explicit argument for the latter claim is to be presented in the subsequent section.

6.3 On the Unreasonable Effectiveness of Semantics

It might be tempting to view the current proposal as a suggestion that formal semantics is an exercise in formalisation or mere formal description. Whereas the theories discussed section 6.2.2. are explanatory in nature, scientific or otherwise, formal semantics merely describes the kinds of compositional facts stated in section 6.2.1. This is the picture against which Szabó (2015) argues.

The working semanticist (who tends to be employed in a linguistics department) is in the business of building what various authors have dubbed a *descriptive* semantics; the speculative semanticist (a philosopher, no doubt) is in the business of thinking about the *foundational* questions of meaning (1).

The inimical idea argued against here is that (descriptive) semantics sets itself the task of simply collecting systematic data, such as the compositional semantic facts, and cataloguing this data for use in future theories of meaning. Szabó argues further that cataloguing can in no way be construed as theory building as “a catalogue remains silent about the reasons behind the classification it employs” (2015, 2). This objection is redolent of similar critiques of the pre-Chomskyan American Structuralist movement in which figures such as Charles Hockett claimed that linguistics is ultimately a classificatory science. This interpretation of the descriptive project is supposedly at odds with a *genuine* explanatory project in which semanticists take themselves to be engaged. I think that there is a way of avoiding these particular problems.

By appreciating the role of mathematical modelling in the sciences, we can find an intermediary position which neither neglects explanation in semantics nor Lewis’ caution against confusing formal semantic models with theoretical linguistic persuasions.

I distinguish two topics: first, the description of possible languages or grammars as abstract semantic systems whereby symbols are associated with aspects of the world; and, second, the description of the psychological and sociological facts whereby a particular one of these abstract semantic systems is the one used by a person or population. Only confusion comes of mixing these two topics (Lewis, 1970: 19).

Models can explain things, they can make predictions and they can even correct theories (as was the case with Fitzgerald's pulley model). In some cases, as in certain physical applications, they can even act as experiments which test the validity of theoretical claims (this is how I have always considered intuition-based data in generative linguistics). It would be a mistake to equate the role of a model with a taxonomy or list of items.²⁰ The explanatory nature of models can be brought out by appreciating the features postulated of them which are meant to capture features of the target domain. For instance, in the case of formal semantics, compositionality is motivated as a tool used to capture various limitations on human cognitive capacities. Szabó states that "there are good explanations in the sciences [...] and one of the things that makes them good is that they track objective dependencies in the world" (2015: 7). I agree with this. But this is precisely the way in which most models are set up, either by means of resemblance relations (Giere 1988) or morphisms, isomorphisms (van Fraassen 1980) or more recently monomorphisms (Bueno and Colyvan 2011) etc.

The claim I wish to put forward is that understanding the explanatory nature of semantic models involves essentially the same puzzle as understanding the success of mathematical modelling within the empirical sciences more generally. I will first outline the philosophical problem of applied mathematics then offer an overview of some contemporary answers which I believe can be applied *mutatis mutandis* to the explanatory role of formal semantics in linguistic theory.

We start with a problem initially presented by the physicist Eugene Wigner (1960). There are essentially two related problems.

1. "The first point is that the enormous usefulness of mathematics in the natural sciences is something bordering on the mysterious and that there is no rational explanation for it"
2. "Second, it is just this uncanny usefulness of mathematical concepts that raises the question of the uniqueness of our physical theories"

Another way to look at the puzzle is that mathematicians do not seem to be concerned with empirical application in their daily work, they are usually interpreted as being motivated by *a priori* considerations, and yet the theories they create (or discover) have remarkable relevance for empirical research. Surprisingly, philosophers of science (and equally, philosophers of mathematics) have had little interest in this puzzle. More recently, however, interest in the topic has been awakened. The

²⁰I think that lists get a bad reputation. In some cases they can prove to be genuinely explanatory. For instance, consider Katz and Fodor's (1963) account of lexical semantics in terms of semantic markers or lists of structures with positive or negative values. This simple but elegant theory can be used to explain various entailment relations between words. Contemporary theta-theory incorporates many aspects of the older view.

most intuitive approach to solving the problem of applied mathematics is the mapping account of Pincock (2004a, 2007).

The mapping account draws inspiration from mathematical structuralism (Hellman 1989, Resnik 1997, Shapiro 1997). On this view, mathematics provides abstract structures which empirical scientists match up to various empirical structures.²¹ By some miracle, the mathematical structures developed by mathematicians yield explanations and predictions in the natural world. “The idea is that there is some structure-preserving mapping between the world and the mathematical structure in question, and that is pretty much the end of it” (Bueno and Colyvan, 2011: 347). In the same way that I could use a map of Edinburgh to plot my course from the train station to the airport, a scientist can use mathematical structures of say group theory or differential calculus to discover facts about the structure of the physical phenomenon under study. Take the case of the Lotka-Volterra model of predation. The physical phenomenon in need of modelling was something which puzzled biologists after the first world war. They noticed that there was a shortage of sea life in the Adriatic despite the fact that fishing had significantly decreased at the time of the war. On the other hand, the number of predatory species had increased dramatically.

Not only did the purely calculus based Lotka-Volterra model shed light on the phenomenon but it also predicted that “the predator and prey populations will oscillate indefinitely, out of phase with one another” (Weisberg, 2013: 12). The equations (below) which determine the mathematical model in this case have biological consequences when equilibrium values are solved for (by setting the equations to zero and assigning values to the variables).

- $\frac{dV}{dt} = rV - (aV)P$
- $\frac{dP}{dt} = b(aV)P - mP$

The above set of formulas, where P and V stand for the population sizes of the predators and prey respectively, imply what Weisberg and Reisman (2008) call the “Volterra Property”. This property states that any event which has a harmful effect on both P and V will result in an increase of the size of V , as the post-world war I biologists noticed. Conversely, events such as light fishing, as occurred during the war, would result in an increase in P or predator population. Puzzle solved!

Returning to the mapping account, the equations set up certain structural relations between objects, these relations can in turn be mapped onto the structures of the physical environment, in the case above predator-prey sizes. By manipulating the

²¹Technically, scientists assume a structure for the natural world, replete with various abstractions and distortions.

mathematical structures posited by the equations we gain insights into the physical phenomenon or the relationships between the physical objects in the target system.

The process can be seen in semantics as well. Consider the puzzle of structural ambiguity. Recall that in section 6.2.1., part of the semantic explananda or semantic facts in need of explanation were sentences such as (2).

(2) Every boy likes a toy.

The above sentence has two readings, one in which each boy (in the domain) likes some distinct toy and one in which each boy likes the same toy. Structural ambiguities like these are ubiquitous in natural language. There are two ways in which to account for this phenomenon in linguistic theory.²²

We can follow Montague (1970) and alter the syntactic model thereby producing two distinct semantic derivations *via* compositionality.

The principle of compositionality requires that every (non-lexical) semantic ambiguity corresponds to a derivational ambiguity. Whenever a sentence has more than one meaning, there should be more than one way of constructing it (GAMUT, 1991).

In Montague Grammar, the alternative readings are generated by “rules of quantification” which is another method of sentence construction. Contemporary syntactic models, however, have more reasonable methods for deriving the requisite structures. Nevertheless, representing the scope ambiguities through alternative syntactic derivations can lead to an explosion of ambiguity and the need for innumerable alternative syntactic configurations. Motivation for syntactic alteration is often *ad hoc* and not well-motivated as was the case in Montague’s original account.

Underspecification is another way to go. Importantly, the underspecification approach does not automatically entail a departure from the principle of compositionality. Semantic underspecification is basically an intentional omission of linguistic information from semantic description. The underlying idea is to postpone semantic analysis until it can be executed in such a way that various ambiguities are resolved. As we have seen in the previous chapter.

The key idea of underspecification is to devise a formalism which allows to represent all logical readings of a sentence in a single compact structure. Such a formalism allows one to preserve compositionality without artfully casting pure semantic ambiguities into syntactic ones [...] (Lesmo and Robaldo, 2006: 550).

²²The two logical forms of the sentences are:

- (1) $\forall x(Boy(x) \rightarrow \exists y(Toy(y) \wedge likes(y)(x)))$
- (2) $\exists y(Toy(y) \wedge \forall x(Boy(x) \rightarrow likes(y)(x)))$

Again, this process amounts to a type of storage of interpretations without immediately checking for consistency or a level of representation which allows for the ambiguities to be unresolved. At a later stage these interpretations are pulled out or extracted and interpreted in parallel. Whether we use Cooper Storage, or holes and labels as in Bos (1995), the formal semantic apparatus can represent ambiguity without complicating the syntax.

Therefore, structural ambiguity can be explained in two ways, by altering the combinatorics of the syntactic models or by accounting for it purely in terms of semantics as in underspecification. The mathematics involved in both approaches is supposed to do more than just represent scope ambiguity in natural language syntax but it is meant to explain it. Why else would there be two non-equivalent approaches, if one did not do a better job of explaining the phenomenon than the other? This is not to suggest that one approach is the *right* approach and the other is not but rather that explanation can be a deciding factor between rival models (where simplicity, elegance and parsimony might be in the service of such a goal).

The idea in the case of the Lotka-Volterra model as well as the scope ambiguity case is that the underlying mathematics can explain the phenomenon in question. Glanzberg (2014) seems to root the effectiveness of semantics in its application of mathematics similarly. “Good explanations tend to appear where we apply model theory or other branches of mathematics to semantics, while mere disquotation signals explanatory weakness” (Glanzberg, 2014: 268). Therefore, it is when mathematics is applied that semantic analyses become explanatory. *Why* is a particular sentence such as (2) above ambiguous? *Because* either it has two distinct syntactic configurations and thus two corresponding semantic representations or *because* certain syntactic objects involve complexities which require additional semantic resources for their interpretation, i.e. interpretation is not immediate and can be resolved in distinct ways.

The mapping account described above is meant to provide insight into how exactly these mathematical structures explain and predict physical phenomena. There are two missing pieces of this account. The first is an account of what the precise structure-preserving mapping is between the model’s mathematical structure and the structure of the physical system, i.e. the type of morphism or resemblance relation. This omission needn’t concern us too much though and can even be approached on a case by case basis. The second part of the picture concerns the “assumed structure” of the target system. This is an unfortunate convenience but a necessary one. Without the assumption of structure, the mapping account (and most others) cannot get off the ground. But the problem of applied mathematics remains. Why should a physical system be beholden to the strictures of mathematical equations, proofs and theories or the structures which they posit?

The main problem for the mapping account is that maps do not really explain anything but merely represent or depict structures and relations between the model and the target.

The problem is simply that it is hard to see how a mere representational system can provide explanations and yet that is the only role mathematics is allowed to play in the mapping account (Bueno and Colyvan, 2011: 351).

The weak (if any) explanatory force of disquotational semantic accounts, discussed by Glanzberg, can be explained in these terms likewise. All that disquotational schemata do is provide mappings from sentences in the object language to corresponding sentences in the metalanguage.

Thus, defenders of the inferentialist conception of applied mathematics expatiate upon the mapping account to include a level of inferential relations. In other words, the strategy of applying mathematics, as in formal semantics, involves establishing “*inferential relations* between empirical phenomena and mathematical structures, or among mathematical structures themselves” (Bueno and Colyvan, 2011: 352). Unlike the mapping account, explanation, novel prediction and idealisation can be accounted for if we conceive of the relationship between models and their targets as an inferential one. First, we build models using mathematics, then we introduce mappings between the models and the assumed structure of the target system, finally we draw inferences from the dynamics of the models (in pure mathematical terms) to the behaviour of the target systems interpreted accordingly.²³ This is somewhat different from standard modelling in that we are not only attempting to model the behaviour of the target system but impose a certain structure onto it for the sake of drawing empirical conclusions from mathematical ones.

Formal semantic explanations and predictions can now be understood in terms of this framework. The explanations involve initially converting the natural language phenomenon into systematic representations of their form through syntactic models. These models or grammars can be further represented as abstract algebras which are homomorphically mapped onto semantic algebras. Ignoring this complication for now, the model-theoretic apparatus (involving types and lambdas as shown in section 6.2.3) structurally mirrors the target system, in this case the syntactic structure (and *not* the empirical setup). The semanticist then draws conclusions from the semantics such as what two interpretations of a structurally ambiguous sentence might look like or what the entailment relations of a given generalised quantifier are using logic or other tools. The final step is then to interpret the semantic conclusions back into the syntactic formalism (this can sometimes be the inverse of the initial mapping). Hence, the explanatory nature of the process comes in at the stage of interpreting the features of the semantic analysis used to model the initial phenomenon. For example, what *explains* the two different meanings of a quantified sentence like (2) above is that there are two separate semantic representations which are interpreted by deriving

²³There are some specifics involved here, namely three steps which include immersion or the initial mapping, the derivation step which is when theorists draw mathematical conclusions in the model and finally the interpretation step in which the previous conclusions are interpreted in the empirical or assumed structure. For more details, see Bueno and Colyvan (2011).

two distinct syntactic forms. In the case of underspecification, the application alters the nature of the problem and subsequent explanation, from syntax to semantics by means of a store or hole which is interpreted individually at a later stage of semantic processing. This is not unlike an application employed by Galileo to solve a problem in the physics of motion in which he represented the motion of bodies diagrammatically and solved it *via* the principles of geometry.

Galileo's strategy is to take a problem in physics and represent it geometrically. The solution to the problem is then read off from the geometrical representation. In brief, he reaches his answer by changing the question; a problem in kinematics becomes a problem in geometry (Hughes, 1997: S327).

To sum up this section, the application of mathematics in semantics operates in a similar fashion, i.e. involving what Suárez (2004) calls "surrogate reasoning". We first formalise the target using syntactic modelling, then using model theory and other tools of mathematical logic map this result into set theoretic objects, then solve various ambiguities, entailments and the like in the set-theoretic structure before attempting to interpret the results back into the initial syntactic framework.²⁴ The account of empirical prediction runs similarly. In the next section, I investigate how formal semantic models can be used to predict the behaviour of linguistic phenomenon through the case of NPI licensing.

6.4 Predicting Syntactic Felicity in the Real World

The previous section set up a connection between formal semantics and applied mathematics. It outlined how modelling generally works in the applied sciences and how models can be genuinely explanatory in semantics. In this section, I plan to expand on the puzzle of applied mathematics and show that one of its sub-puzzles has an analogue in semantics as well, namely the problem of prediction.

It is one thing to argue that applying mathematics to an empirical undertaking can yield genuine explanation based on either mappings or inferences or both. It is quite another to claim that empirical predictions can be made purely by means of consulting the mathematical formalism. In the previous case of the Lotka-Volterra model and the structural ambiguity problem, we used the mathematics, differential calculus and formal logic respectively, to explain the behaviour of the physical systems under investigation.

Unlike the original statement of the puzzle or "unreasonable effectiveness of mathematics" courtesy of Wigner (1960), Mark Steiner (1998) argues that the wonder

²⁴It should be noted that Hughes is not arguing for an inferentialist account in the quote above but rather what he calls DDI or Denotation, Demonstration and Interpretation. However, the two accounts do have a lot in common. Nevertheless, it is not my intention in the present work to take any specific stance on the philosophy of applied mathematics.

of applied mathematics involves not one puzzle but a family of puzzles. Colyvan (2001) considers the puzzle of prediction to be particularly intriguing, “where the mathematics seems to be playing an active role in the discovery of the correct theory” (267). A canonical case of this in the natural sciences is how Maxwell’s equations predict electromagnetic radiation (this is Colyvan’s test case).

The case I would like to present in order to display the predictive powers of semantic models is the modelling of NPI licensing environments. There has been a wealth of linguistic literature on the topic of Negative Polarity Items (NPIs) in the past few decades. Part of the reason that these particles have been so widely studied is that they exhibit strange hybrid characteristics. When they are unlicensed in a particular context (cannot be appropriately used or generally not found) we are left with an infelicity akin to violations of syntactic rules. However, the linguistic situations in which these items are licensed are generally explained purely in terms of semantics (Rothschild, 2009). This is interesting for a number of reasons. For one thing, the NPI phenomenon was never considered to be semantic data before certain entailment relations were noticed with relation to it. Thus, we saw a shift in what counts as data for semantic theory as discussed in section 6.2.1.

A precise definition of NPIs is hard to find since these particles are usually found in various contexts and do not always contribute to the meaning of those contexts. However, according to Hoeksema, “[p]olarity items are words or idioms which appear in negative sentences, but not in their affirmative counterparts, or in questions, but not assertions, or in the protasis of a conditional, but not in the apodasis”(1997).²⁵ Common NPIs are words such as *ever*, *any* (the so-called free choice any will not be considered here), *at all* which appear felicitously in sentences such as (6.10), (6.11) and (6.12) below (but not in their pairs):

(6.10) Xoliswa didn’t *ever* see soccer balls.

*Xoliswa *ever* saw soccer balls.

(6.11) Mandla doesn’t want *any* birds.

*Mandla wants *any* birds.

(6.12) The witness never told them what they asked *at all*.

*The witness told them what they asked *at all*.

The first sentences in these pairs all have something in common, namely they all contain negations. As Ladusaw eloquently put it “NPIs live under the shade of negation”. However, the matter is not this simple. As indicated by Hoeksema’s quote above, there are many other instances of NPI licensing which do not seem to involve negation. Consider (6.13) below:

(6.13) Every person who *ever* visited South Africa, loved it.

²⁵Latin grammar school words for antecedent and consequent of conditionals.

The above example involves the quantifier *every* and seems to license the NPI *ever* without issue. What is needed is a model which not only predicts what negation and quantifiers like *every* have in common but also can explain and predict when NPIs are licensed elsewhere. Standard accounts proceed *via* a notion of downward entailment or DE on predicative contexts. A DE context is one in which you can replace a predicate with a stronger predicate (or more exclusive) without altering the truth of the sentence. All sentences involving negation and universal quantification are DE contexts. Consider the pairs in (6.14) and (6.15) below.

(6.14) Xoli didn't *ever* see soccerballs.

Xoli didn't *ever* see blue soccer balls.

(6.15) Every person who *ever* visited South Africa, loved it.

Every American who *ever* visited South Africa, loved it.

Thus, we can affirm the principle "a predicative context is NPI licensing iff it is DE" in our model (or the Fauconnier-Ladusaw Hypothesis). Job done. Actually, not so fast. DE contexts fail to predict the felicity of NPIs in non-monotonic contexts (neither upward nor downward entailing). Once again, consider the following examples.

(6.16) Most workers enjoyed *any* job which they were offered.

(6.17) Most people who have *ever* been to South Africa, loved it.

These contexts are not downward entailing nor upward entailing since if we substitute a stronger predicate or weaker into (6.16) we are no longer guaranteed of the truth of the sentences.

(6.18) Most workers enjoyed any harmful job which they were offered.

(6.19) Most workers enjoyed any activity which they were offered.

Similarly, adverbs, "only" constructions and the antecedents of conditionals seem to offer counterexamples to previous logical constraint on our models. A more successful approach makes use of tools from model theory to better predict NPI licensing and capture the data so far described. Rothschild (2009) introduces a notion of domain-sensitivity (DS) to model the phenomenon. The logical arsenal necessary is that of a model which is classically defined as containing a set of individuals (and individual events) and an interpretation function which maps predicates and names in the object language to those individuals and sets, $M = \langle D, I \rangle$. "We'll idealize and suppose that a given sentence *S* is either true or false relative to any model *M*" (Rothschild, 2009: 14). The informal definition of domain-sensitivity applies to predicates if a sentence is true in a model then adding more objects from the domain which satisfy the predicate can make it false in that model.

Formally, DS is defined in terms of conservative extensions of the domain:

Conservative Extension – A model M' is a conservative extension of a model M if M' contains all the individuals and events in M , and at least one more individual or event not in M , and the predicate-extensions in M' are the same as in M in so far as they apply to entities of M alone.

Domain Sensitivity itself is then defined as:

Domain-Sensitivity – Given a predicative context c , with a 1-place predicate P in it, a sentence S is a domain-sensitive context if and only if for every model M such that S is true in M there exists a conservative extension of M , M' , s.t.

1. S is false in M' .
2. For each $i \in [M' - M]$, Pi (Rothschild, 2009: 15).

Without too much toil, we can see that this analysis accounts and predicts all of the above occurrences of NPIs (and more). Thus, by applying model theoretic techniques to the syntactic phenomenon of NPI licensing, our semantic models are able to predict a range of data previously unaccounted for.²⁶ Here as before, it is the mathematics doing the predictive work, a characteristic of applied mathematics. But again, mere mappings do not seem to be enough. It isn't enough to say that English and set theory or model theory have a common structure (or structural overlap) but we also need to explain why the inferences from the formal domain are relevant to the natural one. In fact, there are cases in the history of science in which aspects of the model thought to be artefactual proved to have empirical significance. For instance, as reported by Bueno and Colyvan (2011: 364), multiple revisions, in terms of physical interpretations, of the same mathematical formalism in classical mechanics led to the discovery of the positron. Dirac initially thought negative energy solutions was merely features of the mathematical model and not physically realised but later, after finding physical interpretations of these solutions, it caused him to revise his entire theory and predict the existence of a novel particle. In general the mathematical structures applied scientists use are much richer than the physical structures being modelled and this can lead to predictions based on logical extensions of the mathematics or merely interpreting "unused" mathematical structure.²⁷ Importantly, although the literature on NPIs is sometimes presented as a project in falsification, the progress made in this area is not progress by counterexample. Rather, as Williamson (2016) puts it (in a different context), "[w]hat defeats a model is not a counterexample but a *better model*, one that retains its predecessor's successes while adding some more of its own". I cannot think of a better way of describing the research paradigm in NPI licensing.

²⁶Nothing I have claimed so far rests of the definitiveness of Rothschild's analysis. If another model does a better job at the end of the day, the point remains the same.

²⁷Sometimes the opposite is also the case, where a modeller chooses an impoverished mathematical model to isolate a specific aspect of the target system much richer in reality. The Ising model of ferromagnetism in statistical mechanics is an example of this. See section 4.3.

Solving the puzzle of prediction is not my task in the present work but rather my claim is that this puzzle for applied mathematics is shared by its formal semantic counterpart. The central claim remains, in the same way that an engineer uses and applies mathematics as diverse as algebra, trigonometry and calculus to physical problem solving tasks, the semanticist uses model theory, lambda calculus and formal logic to interpret syntactic formalisms or models for the task of explaining and predicting natural language data as in the case of NPI licensing.

6.5 The Syntax-Semantics Interface

So far, we have implied a lot about the relationship between syntax and semantics. I suggested a strong analogy with formal semantics and applied mathematics. I argued that semantic modelling involves setting up systematic mappings between the syntactic formalism or structure and semantic model and then drawing purely mathematical conclusions which are to be translated or read back into the syntactic models under study. One thing that remains to be shown is how exactly to account for what is known in the literature as the syntax-semantics interface. I mentioned before that the precise relationship between model and target can be determined on a case by case basis. The present case is in need of such precisification.

Before addressing the issue of the syntax-semantics interface, there might be one preliminary problem the solution of which can shed light on the overarching goal of this section. The details of applied mathematics discussed in the previous two sections suggest that mathematics is somehow effectively applied to an empirical or real world target system. I have attempted no such thing in my exposition of formal semantic modelling. Now, as previously mentioned, most accounts of applied mathematics (especially mapping accounts) are forced to impose a structure on the empirical domain. But this is different. I was not assuming that generative syntax (or categorial grammar or whatever) is the “assumed structure” of the target system. In fact, I do not believe that it is. Syntax involves many distortions and idealisations which simplify and ignore linguistically relevant aspects of natural language (see Part II.4 of this thesis). In addition, natural languages also involve phonological and morphological structure. The key to appreciating the analysis of formal semantics as applied mathematics presented here is that applying mathematical models (much like syntax) is a recursive endeavour. The procedure I outlined above, of mapping then derivation then interpretation (reverse mapping sometimes), can be applied to other mathematical models taken as the target or empirical setup which in turn can undergo the procedure again. In effect, formal semantics, as I view it, is the process of applying mathematics to other mathematical models, namely syntactic models.

We can see this when we consider the varieties of semantic models offered on the basis of different syntactic formalisms. Montague himself modelled his semantics on a categorial grammar formalism but the initial syntactic rules used in derivations

were critiqued for being *ad hoc* and purely driven by the semantics. A more sober categorial syntax-based semantic model is offered by Jacobson (2014) within her direct compositionality programme. The idea behind direct compositionality is that unlike standard approaches to semantics (such as found in Heim and Kratzer (1995)), semantic interpretation is not merely epiphenomenal or some sort of a logical form (LF) of an otherwise autonomous syntax.

The hypothesis of Direct Compositionality is a simple one: the two systems work in tandem. Each expression that is proven well-formed in the syntax is assigned a meaning by the semantics, and the syntactic rules or principles which prove an expression as well-formed are paired with the semantics which assign the expression a meaning [...] It is not only the case that every well-formed sentence has a meaning, but also each local expression (“constituent”) within the sentence that the syntax defines as well-formed has a meaning (Jacobson, 2014: 9).

There is a strong rule-to-rule mapping between the syntactic formalism and the semantic model on this view. In light of this, categorial grammar seems to work well since the formalism posits functional syntactic composition rules which seem to match up nicely to the functional type constructions of the semantics. Jacobson of course insists that the use of categorial grammar is a mere convenience and her project is compatible with various syntactic formalisms. I do not doubt this, but there are features of categorial grammar that seem to make it more amenable to the kind of mathematical structures used in formal semantics (see section 6.2.3.). There is at least a choice point here. Either we model syntactic phenomena with something more akin to generative syntax (along the lines of Government and Binding theory (1981)) or we model natural language syntax with later semantic modelling in mind as in the case of formalisms such as categorial grammar. The choice is not trivial but it is a modelling choice which can take a number of considerations in mind, e.g. computational complexity, fidelity to the target system, even machine learning or translation applicability.

The point is that models can be applied to other models *ad libitum*. The behaviour of the empirical system can sometimes be buried under heaps of applications each one modelling the dynamics of the previous formalism with the hope that the dynamics of the models percolate up the process from the underlying empirical structure and the target system is not completely lost.²⁸ The picture I am presenting might be obscured

²⁸This picture resembles some accounts of serial, as opposed to parallel, processing in computer science. Given the division of labour in most science and linguistics department, I do not think that this is an unrealistic depiction of the situation. However, there might be exceptions to the current view of mathematical application in linguistics. Perhaps, Jackendoff’s parallel architecture constitutes multi-directional or parallel modelling aimed at the same target. I am not sure about this though. See Jackendoff (2002) and Jackendoff (2007) for an overview of his approach to generative grammar. Even if we do not take this option, the apparent temporal order of mathematical application is a feature of the exposition not the phenomenon itself.

somewhat by the fact that many semanticists make use of very impoverished syntactic representations for the purpose of their analysis. This sort of abstraction, however, is common practice in scientific modelling. The use of simple trees (of which I myself am culpable in section 6.2.3) is merely an explanatory convenience and when a more detailed semantic model is called for, representing the syntactic structures more accurately becomes pertinent.²⁹

The point about the recursive nature of mathematical application can perhaps best be appreciated by considering cases within mathematics itself. It turns out that many areas of mathematical investigation which are considered “pure” mathematics are often applied in the sense I have been stressing above. Consider Hilbert’s early work on relative consistency proofs. For instance, in *Grundlagen der Geometrie* [1899], he attempted to show that geometry could be reduced to real analysis and the consistency of the latter is enough to ensure the consistency of the former. Thus, real analysis was applied to geometry and features of analysis (such as its consistency) were then meant to be interpreted or translated back into geometry. The entire Logicist programme of Frege, Russell and Whitehead can be seen as an exercise in the application of formal logic to mathematics. Granted, reduction and modelling are not the same thing. Nevertheless, contemporary mathematicians are much less interested in the foundational questions of the early 20th century (or the search for a one true stable foundation for all of mathematics), yet they still apply various mathematical theories to one another with fruitful results, such as algebra to topology, set theory to arithmetic and so on. Not all theories under the banner of pure mathematics aims at describing mathematical reality (whatever that is) directly.

We do not have to go to mathematics to find examples of the embedded modelling paradigm. Morrison (2015) offers an example from population genetics which utilised vast amounts of mathematical apparatus in modelling real world populations so much so that a similar characterisation (as the one I am offering for semantics) is appropriate in this domain.

Instead of the natural or “real” populations studied by field biologists, populations are now often mathematical constructs that can be manipulated using sophisticated mathematics. In this context, the model, rather than the real-world environment or system (that is, populations of living things), becomes the object of inquiry (Morrison, 2015: 23).

Returning to the issue of the precise relationship or morphism between the syntactic and semantic models, there are a number of worries. For one thing, isomorphisms and even homomorphisms seem to impose too strong a constraint. To see this, let

²⁹Even very different semantic formalisms such as Discourse Representation Theory (DRT), which is intended to be interpreted cognitively, presupposes a complete syntactic model. This is brought out by one of the chief problems with DRT, namely that it has been argued not to capture the incremental nature of semantic parsing (hence the move to frameworks such as Segmented Discourse Representation Theory). I thank Hans Kamp for pointing this out to me.

us consider a related problem for the mapping account in the philosophy of applied mathematics, the problem of idealisations. As we have seen in Part I.3.1 and Part II.4, scientists often use non-veridical models to reflect features of the target system at hand. In other words, models are not always meant to truthfully represent the target system. They involve many distortions and purposes other than truthful representation. What is Fisher's three sex biological model "mapping" in the real world, or natural language as denumerably infinite in Chomsky? We have surveyed some of these models in the previous subpart and part I. The point is that not all of the features of the model can receive realisations in the physical target system as a strong morphism would require.

Similarly, strong morphisms between syntactic and semantic structures seem too strong. Syntacticians posit many features which do not seem to make it to a level of semantic interpretation. For example, obligatory subjects in English as in sentences such as *It is raining* are usually explained by means of what is known as the extended projection principle (EPP). The principle states that subjects are mandatory in *DP* or *NP* clauses even when there is no semantic subject or agent in the surface form. In addition, some syntactic objects are covert as in *PRO* and traces and other members of the empty categories or the category of linguistic objects with null phonology. Take *PRO* for instance. *PRO* is a pronominal object postulated in the subject position of non-finite clauses as in 6.20.

(6.20) Noluntu_i wants PRO_i to meet up.

Szabó (2015) argues, in accordance with Frege's context principle, that empty category words or items cannot be semantically interpreted on an individual basis. These *sub rosa* syntactic elements are context-sensitive and emergent phenomena.

This is exactly why it could not be a semantic primitive: just like the meanings of 'the' and 'is', the meaning of *PRO* is not the sort of thing we could explain without relying on other expressions with which it is in construction (Szabó, 2015: 27).

If some syntactic elements cannot receive individual semantic treatment, then we have a problem for accounts such as direct compositionality and any strict morphism based mapping in general. Furthermore, semantically null elements such as those posited by EPP militate against the idea of a complete mappings from syntax to semantics as strongly if not more.³⁰ There are a number of possible solutions to this problem, the inferentialist conception of Bueno and Colyvan (2011) for applied

³⁰On the other side, Stanley and Szabó (2000) argue from semantic theory to hidden elements in the syntax in order to explain quantifier domain restriction. On this view, *NPs* have covert argument places for domain restriction in contexts. This is an excellent example of semantic modelling being interpreted back into the syntactic formalism as discussed in the previous section.

mathematics and Suárez (2004) for scientific representation in general are two.³¹ Compositionality itself might just be an inferential relation on this view. Another approach could be the partial structures framework of Bueno, French and Ladyman (2002), in which partial homomorphisms allow for some elements of the the relevant domains to remain undefined. These options, however, are not available to a linguist who does not appreciate the applied mathematical nature of her enterprise.

The first step to recovery is admitting that you have a problem. In this case, the problem of the syntax and semantics interface can be resolved by appreciating that syntactic modelling involving empty categories and other covert operators are essentially scientific idealisations. This explains why they pose a problem for mapping-based accounts of the interface such as direct compositionality.³² And only when we admit that these idealisations are in need of explanation at the level of the interface, can we make use of inferentialist or partial mapping accounts of the phenomenon.

Accepting that covert syntax is idealisation can go a step further in resolving some philosophical disputes concerning the ontology of natural language. An alleged problem for Lewis' conception of natural language, reignited by the recent heated correspondence between Devitt (2006, 2008b) and Collins (2007, 2008c), is that the notion of convention and (nominalism in general) cannot explain various unvoiced syntactic elements such as *PRO*. What corresponds to *PRO* in a linguistic community's convention establishing behaviour? Chomskyan mentalism is usually proffered as the only theory capable of accounting for these postulates since it assumes a level of syntactic representation beneath the surface syntax. External languages, on the other hand, have recourse only to overt elements of speech and communication.

From the modelling perspective advanced here there is nothing difficult to explain. Idealisations in the syntax can be approached from various angles as convenient distortions or falsehoods introduced for the purpose of tractability of the target system whether this is an external language or an internal mental one. The semantic models are then applied to the formalism or scientific model used to model syntactic phenomena. The syntax-semantics interface is accounted for purely in terms of the application of one mathematical formalism to another which is in turn mapped onto the assumed structure of the empirical target system. These mappings are generally partial in nature allowing for elements in either or both relata to be undefined for some elements (the same can be applied to phonological and pragmatic models at the interfaces). This picture dovetails with Glanzberg's recent analysis of semantics. "When it comes to explanatory force, our semantic theories are for the most part partial. Where they rely on mathematics, they are explanatorily substantial, but where they rely on disquotation they fail to be" (Glanzberg, 2014: 277).³³

³¹Interestingly, Brandomian inferentialism, in its rejection of the denotational truth-conditional semantic orthodoxy, would be similarly immune to these context-sensitivity worries.

³²Jacobson does address these elements and offer a way of capturing them within her framework.

³³His account does, however, rely heavily on the posit of linguistic competence and the faculty of

6.5.1 Semantics without Semantic Value

Earlier, I discussed Lewis' (1975) view in the context of it being an intermediate approach compatible with various takes on the ontology of natural language (or target of linguistic models). There is, however, a stronger reading of Lewis on which mentalism (and subsentential semantics) does not find a natural home as we saw in Part I.3.1. For Lewis, a language can be captured by a list of pairs of forms and meanings. The notion of language or the function which pairs sentences with their intensions (the function accessed by linguistic communities) is not very fine-grained. The semantic values which our models describe, on this view, are at the sentential level which might be considered to distance them from the compositional semantic data of section 6.2.1. Yalcin describes the point in the following way.

Lewis expressed no interest in accounting for productivity. On the contrary, he explicitly expressed skepticism that semantics could be supplied with determinate foundations at the level of subsentential expressions. He offered only to ground the notion of a population's using a language, where 'language' in his sense refers just to a function which pairs sentences with meanings, not the far richer object assumed above. This limited set of facts of 'semantic value' were to be grounded in certain conventions—on Lewis's analysis, in certain regularities in belief and action prevailing in a given population, owing to some common interest in communication (Yalcin, 2014: 39).

This discussion brings in an important aspect of the picture I have been pushing thus far. The concept of semantic value of an expression or subsentential expression is of core interest to any account of semantic modelling. On my view, semantic values are offshoots of syntactic wellformedness, thus they do have subsentential significance *contra* Lewis.³⁴ However, the embedded applied mathematical framework depicted here, takes no stance on whether there is anything in the world of natural languages (or empirical setup) which directly corresponds to these semantic values. They could be artifacts of the models or emergent phenomena at the level of interface or many other options. I argued that semantic modelling is indirect (perhaps multiply so) and the functions, types, lambda formulas used to describe the models have no direct bearing on the target system of linguistic models generally, namely natural language phenomena, but rather bear relations to the more circumscribed target of syntactic models. Typed lambda terms map onto syntactic constituents which in turn model utterances or meaningful discourse (even the notion of a 'sentence' is theoretical to an extent).

language as being the ultimate goal of a linguistics and linguistic semantics.

³⁴In other words, Lewis might have been taking semantic models to be applied indirectly to the linguistic target, i.e. conventions of communities. However, on my view, this is not the whole picture but rather semantic models are applied to syntactic models (which are usually hierarchical and productive in Yalcin's sense). At the end of the day, we might still be in the business of accounting for linguistic regularities in various communities but *via* a far more indirect methodology.

This leaves the question of the ontological status of semantic values open. On my view, there could be something in nature which indirectly corresponds to semantic value and percolates up from the syntactic formalisms to the level of semantic modelling. Of course, the reverse could also be true and there could be no natural phenomenon at the base of the embedded modelling, or at least no natural individual feature beyond the general phenomenon of “meaningful discourse”. Rayo (2013) embraces such a paradigm by endorsing both *Semantic localism* and a “Grab Bag Model” of linguistic competence. The former is a position which assesses the acceptability of an assertion purely in terms of local context or more specifically “a localist would claim that all that is required for an assertion to be in good order is for it to succeed in dividing *the possibilities that are relevant for the purposes of the assertion* into verifiers and falsifiers” (Rayo, 2013: 650). The Grab Bag Model, as opposed to the “Specialised-Knowledge Model” of generative linguistics, states that there is no independent module comprised of general semantic rules but rather a hodgepodge of mental entities such as memories, images, dictionary entries, personal anecdotes etc. associated with lexical items. Importantly for our purposes, Rayo holds that the correct picture of semantic interpretation is somewhat as follows.

It would be better to imagine that your companion hands you the grab bag plus a bit of ‘syntax’: an explanation of how the grab bags should be combined to render one of the relevant possibilities salient (2013: 649).

Syntax still has a role to play in semantic interpretation. The individual grab bags different people might possess could contain distinct mental entries for the same lexical item but I would opine that the syntactic combination principles would have to remain rather constant among individuals in a given linguistic community in order for their grab bags to be useful at all. In other words, we could allow for some variation among language users in terms of lexical items and their meanings but if we all meaningfully combine these items differently we would be at a linguistic loss. Rayo’s view is similar to other usage based accounts which take communication to be the primary linguistic explanandum as well as motivation but he is forced to admit that in order to be understood, your interlocutor “hands you some syntactic information” (2013: 649). Thus, the model is still that of given some syntactic information you are expected to attach semantic significance, perhaps just more incrementally.

One difference between this view and traditional semantics is that semantic significance is not a global or complete set of semantic rules but a motley assortment of cognitive and contextual information based on relevance to particular situations. It is, therefore, unsurprisingly that Rayo eschews of traditional linguistic meaning in favour of a more Stalnakerian account in terms of contextual tracking and updates as well as general cooperation principles (such as his *Principle of Clarity* (654)). Neither compositional semantic facts, as I have described them, nor truth-conditions seem to receive explicit treatment on this view. However, as I have stated, on the

view of formal semantics as applied mathematics there is no direct commitment to individual semantic values or linguistic meaning *as such*. Rayo's rejection of invariant linguistic meaning is compatible with the model of semantics presented thus far. The mathematics he employs to interpret syntactic information is indeed distinct but the target he is aiming to explain need not be so radically different. In order to appreciate this point, I will mention three related types of targetless modelling strategies discussed in Weisberg (2013), namely *generalised modelling*, *hypothetical modelling* and actual *targetless modelling*.

One type of investigation involves the construction of models in order to study general phenomena such as parasitism or sexual reproduction. A second is when theorists construct models to study nonexistent phenomena. A third type of investigation involves studying a model with no target at all (2013: 114).

The first kind of modelling occurs when the target of the model is a generalised not specific one. This would be like trying to explain natural language *simpliciter* without isolation into specific modules or subfields (Clark (1996) has a view along these lines). The second modelling strategies involves collapsing the lines between real and unreal systems such that knowledge of the former can be gained by insights into the latter. Fisher's three-sex model of sexual reproduction falls into this category. One could imagine a linguist similarly constructing "impossible grammars" in order to test a particular framework. Some work has in fact been conducted on this topic in order to illuminate UG under the principles and parameters banner.³⁵ Lastly, the final type of modelling is related in many ways to pure mathematics. When theorists engage in this sort of practice, they tend to describe the models for their own sake or for the sake of mathematical manipulation.³⁶ Weisberg cites simulations associated with the field of cellular automata as being studied "perhaps for the sake of what they tell us more broadly about computation" as an example of this targetless category.

It is easy to see that Rayo's proposal does not fit with the last subcategory of targetless modelling since there is a clear target of his approach. However, the target is not a specific system such as the compositional semantics of the previous sections, rather it is a generalised target or phenomena connected with meaningful discourse generally, i.e. the first type or *generalised modelling*. Indeed, if this were the case, it would mark a departure from the syntactic target directed modelling of the previous sections.³⁷

³⁵"Similarly, knowing something about UG, we can readily design languages that will be unattainable by the language faculty." (Chomsky 1991a: 40)

³⁶Perhaps one way to succour Platonism (or at least methodological Platonism) in linguistics could be to say that it is targetless modelling in the last sense. This is not a line I will pursue here.

³⁷But I think that this is rather the point of Rayo's alternative.

6.5.2 Direct Modelling

On the other end of the spectrum, there is a more direct kind of semantic theorising in the contemporary research, one which places semantic values at the forefront of the target system directly. A recent proposal advocated in Szabó (2015) stands in contrast to the view of semantics as applied mathematics presented above. Szabó addresses the folly of assigning type theoretic semantic values to syntactically defined parts of speech and the idea that grammatical notions intervene in such a way as to prevent grammatical categories receiving purely semantic treatment. The orthodoxy in linguistics aims to define parts of speech such as noun, verb, adjective, adverb etc. in terms of syntactic distribution as opposed to the more intuitive semantic descriptions (such as noun=object, verb=event etc.). In my view, the original division of labour stems from the strong claim of the autonomy of syntax within early generative grammar. Nevertheless, formal semantics then provides a logical translation for these categories in terms of type theory, lambda calculus and so on, as I have argued above. Thus, there appears to be a disconnect between grammar and logic, at least at the lexical level. Szabó, however, thinks that this detour to the semantic value of lexical categories is an unnecessary (and problematic) one.

A few preliminaries are needed in order to appreciate the proposal on offer here. First, the distinction between open and closed classes of lexical categories. Although this distinction is a staple of contemporary linguistics, its precise definition is surprisingly hard to pin down. One way to attempt a characterisation is *via* the notion of language change. For instance, prepositions tend to remain somewhat constant over the lifetime of a given natural language whereas nouns and verbs undergo frequent changes and additions.

Closed categories tend to be small; their members tend to be short. Closed categories are more stable: we easily coin or borrow new nouns or verbs but when it comes to complementizers or inflections, change is slow and gradual. In addition, open categories do but closed categories don't participate in derivational morphology (Szabó, 2015: 7).

The project aims to model the major lexical categories directly in terms of open class constants with the result that the gap between grammar and logic is reduced. So nouns become not types corresponding to distributionally defined syntactic objects but rather open lexical constants used for reference such that the semantic clause only needs to involve a universal quantifier and a variable specified in terms of reference. Verbs, on the other hand, are constants which purport to predicate. For more specific details, see Szabó (2015).

One thing to appreciate is the radical nature of this proposal. Not only is type theory otiose on this view but the semantic modelling is of a direct nature. In a sense, the juxtaposition of Szabó's account with the standard accounts lends credence to the interpretation of the latter modelling paradigm described in the previous sections.

What Szabó is proposing is a direct modelling strategy for semantics, i.e. one that maps semantic models onto the empirical target directly. Semantic values are part of the assumed structure of the world and we, as semanticists, directly represent these values with our theories.

The direct strategy employed within this suggested framework is related in many ways to what Weisberg (2007a) calls *Abstract Direct Representation* or ADR. An example of ADR in the natural sciences is (allegedly) Mendeleev's construction of the periodic table.

Mendeleev decided to focus his attention on finding trends in the properties of valency, isomorphism, and, most importantly, atomic weight, abstracting away from all of the other properties [...] When the elements were properly ordered, Mendeleev argued, one could see the periodic dependence of elemental properties on their atomic weight (Weisberg, 2007a: 212).

The difference between Mendeleev's construction and Volterra's is that the former aims to represent real properties on the target system directly. No separate model or surrogate reasoning *via* a model was involved with Mendeleev's theorising. Furthermore, the periodic table was not just a device for classification but the basis upon which fruitful predictions were made, predictions which led to the discovery of new elements (e.g. gallium, scandium, and germanium). In the same way that chemical elements are *real* features of the world, for theorists such as Szabó, semantic values are *real* features of actual linguistic reality.

Another semantic framework which employs more direct methods of investigation is distributional semantics. The slogan for this kind of semantic analysis is "you shall know a word by the company it keeps" (Firth 1957:11). The idea here is that a comparison of linguistic contexts can yield insights into the meanings of individual linguistic units. Naturally, distributional semantics fits well with corpus linguistics and statistical methods. Once again, semantic values are directly attributed to the linguistic contexts or to be "found in nature". One aspect of this framework is known as the distributional hypothesis presented below.

The degree of semantic similarity between two linguistic expressions A and B is a function of the similarity of the linguistic contexts in which A and B can appear (Lenci, 2008: 3).

Thus, the semantic value of words and expressions is a "collocation" function directly computed on actual linguistic contexts. This picture does not dovetail with traditional formal semantics and the modelling practices which it involves. But in the same way that Szabó's proposal explicitly stands out as a departure from formal semantics with type theory, so too does distributional semantics stand out from standard semantics as applied mathematics.

The tradition that started with Frege and Russell, based on compositionality, formal logic and type theory is a different animal from statistical or distributional approaches to semantics. In addition, it is this tradition against which Szabó's alternative is set. It is certainly the case that more direct theorising methods are possible with relation to semantic explanations and predictions, however, it is my contention that the dominant approach can best be philosophically accounted for by the applied mathematical analogy I have put forward in this chapter.

6.6 Subconclusion

In this chapter, I have attempted to shed light on the nature of modelling in formal semantics. I have argued that semantics is a form of applied mathematics which indirectly models syntactic formalisms by means of model theory and other mathematical tools. I further argued that semantic models correspond to structures which themselves are involved in the modelling process.

In the same way as an engineer might use geometry to draw conclusions and make predictions about physical structures in the world where there are no pure geometric shapes, the accounts I sketched involved mappings between the empirical domain or the target of the model (which might be yet another model), and inferences or interpretation of some sort. Thus, syntactic models can influence semantic predictions and explanations based on their own structures and semantic models do not directly apply to empirical reality whether this reality is semantic knowledge or communal communicative conventions.

The view espoused here is that formal semantics is parasitic on syntactic modelling. This is not necessarily a novel insight. What is novel about my analysis is that it can account for the explanatory and predictive power of formal semantics in a way consonant with contemporary views in both the philosophy of science and the philosophy of applied mathematics. Additionally, I extended the notion of idealisation in scientific modelling to cases of supposed mismatch between syntax and semantics at the interface. I showed that various puzzles concerning covert material, conventions and external or public language can be resolved by the adoption of this perspective. The idealisations of syntactic theory can in turn be accommodated by various proposals involving partial mappings and inferential relations currently on offer in the philosophy of applied mathematics.

Overall, the picture of formal semantics as applied mathematics can be viewed as a renewal and marked development of a philosophical account of linguistics initially proffered by Lewis (1975). Semantics, on this view, genuinely explains and predicts and the problems which have traditionally beset its philosophical underpinnings with relation to these goals are part of a larger goal of explaining the unreasonable effectiveness of mathematics in the empirical sciences.

Chapter 7

Part Conclusion: Structural Realism

The foregoing chapters were written with an explicit focus on methodology. The overarching purpose was to establish a viable modelling paradigm in linguistics. The reasons for this stance were informed by the general discussion in Part I concerning the first grade of mathematical involvement for linguistic grammars. This grade urged a separation between methodology and ontology. The discussion mostly involved various possibilities for the scientific and philosophical interpretation of grammars in terms of formal mathematical models and their appropriateness within the context of linguistic theory.

However, the aforementioned discussion and subsequent chapters left open a few questions about which a philosopher of science might be concerned. These questions pertain to the overarching framework, in terms of the philosophy of science, to which linguistic modelling belongs. Does the view entail a thoroughgoing instrumentalism? Anti-realism? Is the view informed by van Fraassenian constructive empiricism? Should we be realists about mathematical models? I think that a plausible story could be told for all of these possibilities. Nevertheless, I will opt for a form of realism in this short chapter in keeping with both the discussion and resolution of Part I and the forthcoming ontological claims of Part III. In other words, this section will be a suggestion as to the scientific import of the modelling perspective within the larger context of this work. It is not meant to be a decisive nor a particularly in depth examination but it serves rather to suggest a position on these matters which I believe to not only be consonant with the rest of the thesis but in many ways a natural theoretical setting for it. Therefore, the following will be somewhat speculative and exploratory.¹

If nothing else, this chapter serves to offer an alternative to the instrumentalism which might otherwise be suggested by the first grade of mathematical involvement and Part II. This summary note should not be taken as a detailed or meticulous argument in favour of structural realism but merely a consistent suggestion of how

¹The necessity of such a chapter was impressed upon me by the seminar participants of a talk I gave at the School of Philosophy, Religion and History of Science at the University of Leeds in 2015.

to possibly unify the various parts of the present work.

One problem with instrumentalism is that it makes a miracle of the work of models. Modellers are not taking blind shots in the dark and chancing upon useful models. They are informed practitioners guided by their knowledge of formal techniques and sophisticated observations of their target systems. Linguists are no exception to this. Just as we might ask why computational models work for a given task, we can ask why explanatory linguistic models explain anything (if they do)? My simple answer is that it is because they *indirectly* pick out real structures of natural language.

A Structural Alternative

When caught between the pull of realism and the rational scepticism of anti-realism, structural realism has often been considered a happy medium (the “best of both worlds” strategy). I want to sketch some of the reasons for this alternative within the context of linguistic theory in this conclusion before moving on to the philosophical implications of the view in terms of ontology in the final part.

It is well-known that traditional realism in the philosophy of science faces a serious challenge often referred to as pessimistic meta-induction or the problem of radical theory change. This problem relates to explaining progress in science. If our theories are true of the world (or even approximately so), then how can we explain scientific progress in cases in which theories have radically altered (as in the move from Newtonian to Relativistic physics)?

One answer to these sorts of worries is scientific anti-realism. On views under this framework, scientific theories need only be empirically adequate (get the observables right). Van Fraassen (1980) is one case of this view. Interestingly, this latter work has led to much of the focus on modelling in contemporary philosophy of science. Although this might be a viable option, it does lead to similar worries to that of instrumentalism in rendering the success of our models or grammars inexplicable. There is, however, a more modest alternative in views under the banner of structural realism. As Ladyman (1998) puts it,

Rather we should adopt the structural realist emphasis on the mathematical or *structural* content of our theories. Since there is (says Worrall) retention of structure across theory change, structural realism both (a) avoids the force of the pessimistic meta-induction (by not committing us to belief in the theory’s description of the furniture of the world), and (b) does not make the success of science [...] seem miraculous (by committing us to the claim that the theory’s *structure*, over and above its empirical content, describes the world) (410).

Contemporary linguistics faces a similar situation to that of the various paradigm shifts in the history of science. The dominant tradition, generative grammar, is meeting with increased resistance and alternative frameworks such as Dynamic Syntax, HPSG, Construction Grammar abound. Understood in structural realist terms, this

does not entail abandoning many of its insights or successes. Linguistics, like the natural sciences, does not begin *de novo* with every theory change, if we maintain the continuity of *structure*. Seen in this light, Part II.4 and 5 argued for structural relations or similarity between not only different strains of the generative tradition but also across other frameworks such as DS.²

The structures in question are the mathematical models of the theories or the grammars.³ In Weisberg (2013), he describes a third kind of model besides the concrete and mathematical ones, namely computational models. To a certain extent, it is not clear how distinct computational models are from mathematical models (as Weisberg seems to admit when pressed). Nevertheless, computational models have a distinctive procedural or algorithmic element. This aspect allows them to track or represent dynamics of systems (in terms of states and transitions between them). I believe that the models of generative grammar (and dynamic syntax) are of this variety.⁴

The model in DS still involves tree structures and relations on nodes but it extends this analysis beyond the constraints of the generative picture. We have seen (in section 5.2) that the model base in DS also includes multi-person dialogues. However, extending the generative models in line with GB and other frameworks might also be possible and would then cover the same data (be empirically equivalent). Thus, the idealisations would be similar (operations and relations on trees), both would involve procedural computational models and both would cover the same data (as opposed to say the “flat structures” of dependency grammars). The idea, of course, cannot be that the models are identical since they are not (how else would there be progress if there in fact is). But rather the claim was that generative grammar and dynamic syntax make use of similar structures, here conceived of as families of computational models.

²There is precedent for the extension of the structural realist analysis beyond the natural sciences. See Kincaid (2008) for such an account for the social sciences.

³Taking the models themselves to be the structures of a structural realist account is also not unprecedented. On Morgan and Morrison’s (1999) account, models are (partially) independent of theory and the target system as discussed above. Autonomy of models is also argued for in Cartwright and Suárez 2008 with relation to the Londons’ model of superconductivity. In addition, some philosophers of science, such as van Fraassen (1980), take theories to be collections of models.

⁴Here might lie another difference between derivational and non-derivational or model-theoretic grammars, in that the former and not the latter can be considered to be computational (where computational is understood as a proper subset of mathematical models). Chomsky (2000) seems to have something similar in mind when he discusses the difference between the *derivational approach* and the *representational approach*, the former is meant to be understood as a genuinely algorithmic construal of the brain’s actual design *vis-à-vis* generating linguistic expressions and the latter is to be understood as a “direct recursive definition” or conditions on expressionhood (as in the model-theoretic case). Despite claiming that the differences might be overstated or merely intertranslatable, he goes on to adopt a derivational approach under the assumption that it does hold unique insights into language (and additional questions concerning it). The first two chapters of this part can be seen as an account of wherein this difference lies exactly, i.e. modelling strategies.

The above situation is similar to the case of the Londons' model of superconductivity. The previous model was limited in explanatory power and scope. For instance it could not account for the Meissner effect which is the expulsion of magnetic fields from superconductors during the transition to the superconducting state. In order to account for this effect, the Londons took superconductors to be diamagnets as opposed to ferromagnets, a modelling choice independent of theory (or so it is argued by Suárez and Cartwright 2008). In the same way that the Londons' model is claimed to have borrowed piecemeal from other models and theories, DS too borrows from other models, some generative some model-theoretic, and theories, some static and others dynamic, in order to account for anaphoric relations beyond the sentence and person boundary (as well as effects such as those in (1)). The model structure or scaffolding (*via* trees and relations and constraints on subtrees) remains constant. By appreciating the concept and use of models, we gain a clearer picture of theory change and theory comparison which helps to forge a closer tie with the structural realist position in the philosophy of science thereby providing potential answers to the questions of progress and change in linguistics.

This situation is, however, dissimilar from some of the usual conclusions drawn from the recent flurry of formal proofs as to the *weak equivalence* of various grammar formalisms (i.e. they generate or produce the same sets of expressions/sentences). The idea being that Minimalist Syntax (MS), Phrase-Structure grammars (PSG), Tree-substitution grammars (TSG), Head-Driven Phrase Structure grammars (HPSG) or Dependency grammars (DG) are all really just "notational variants" of one another with little empirical consequence (as in Chomsky's (2000) example of $25 = 5^2$ vs $5 = \sqrt{25}$). To a working linguist *qua* modeller, I argue, these proofs mean little to nothing. For instance, dependency grammars posit structural relations which differ significantly from phrase-structure grammars (in fact, DG is flat structurally as opposed to hierarchical, i.e. argument form trumps dominance relations). Similarly, TSG's lack a mechanism for deriving rules such as adverbial modification (easily specified in PSGs) since they do not possess an adjunction operation as in later Tree-adjointing grammars. Yet many of these formalisms can be shown to be weakly equivalent. As Rambow and Joshi themselves note, these equivalences are of little consequence to the syntacticians working in a given syntactic framework who still go about their daily business in very different ways. "The result is a dependency tree, CFGs and TSG are weakly equivalent. However, to a linguist, they look very different" (1997: 3).

Structurally, the models are quite distinct, i.e. they differ in strong generative capacity. By appreciating the roles and operations of the models themselves, we can arrive at a more nuanced account of theory similarity and dissimilarity in linguistics, as I hope to have shown.

Thus we can be realist about the structures indirectly picked out by the models and at the same time be instrumentalist about the models themselves. In the next part I will attempt to find ontological grounding for the structural realism suggested here.

In light of what we have seen in this part, theories cannot just be collections of models, since we have seen that models operate in partial autonomy from theories. Models, however, could still inform theory construction. Cases such as Dirac's discovery of the positron, where an assumed artefact of the model was found to be ontologically significant to the target system, show that it is possible that the structures indirectly represented by models can eventually be described directly by theory.

Part III

**Languages and Other Abstract
Structures**

The dominant picture of the foundations of linguistics and the ontological status of linguistic objects is provided by the conceptualism founded by the generative movement of Chomsky (1965). On this account, languages are mental states, or I-languages, of the individual language users. To 'cognise', or more controversially to *know*, a language is thus to be in a particular cognitive state of the language faculty. This is a physicalist view. All talk of the mind or mental states is just physical talk about the brain at a different level of description (see Chomsky 1986). Hence linguistics is really biolinguistics and is eventually to be subsumed by neuroscience or biology itself.

In the wake of this picture of the foundations of linguistics, linguistic Platonism or realism emerged. Drawing strength from the analogy with mathematics (specifically arithmetic and set theory) and issues within ontology that proved difficult for a physicalist account of the science (in its current form at least), Katz (1981), Katz and Postal (1991) and most recently Postal (2003) offer a radically different account of the objects of linguistic theory and the place of its science. As is the case with many paradigm shifting or challenging notions, the view of linguistics as a formal science did not hold much sway among contemporary practitioners. However, I believe that it does hold genuine insights and approaches the field with bold honesty in interpreting linguistics as *it is* rather than as *we hope it to be*.

In this last part, I will not mount a direct attack against conceptualism (see part 1.3.3 for an argument to that effect). Rather I will take seriously the challenge presented by Platonism while (hopefully) developing a novel account which makes use of some of its core features to better effect. In many cases, I think the words of Katz and Postal lend themselves to my account more so than to naive Platonism (see section 10). Another way of putting my agenda is that I hope to divorce the notion of linguistic realism from that of its Platonist counterpart.⁵ Specifically, the strategy I plan to employ will be to identify three essential desiderata or properties of natural language for which any realist theory of linguistic foundations and ontology ought to account. These properties stem from critiques of the biolinguistic or generative programme offered by Platonists such as Katz and Postal (and nominalists such as Devitt). I will then show that Platonism offers an approach to dealing with these desiderata at too large a cost. Finally, I shall provide a non-eliminative structuralism for the foundations of linguistics in its stead in accordance with a similar interpretation of mathematics (Shapiro 1997, Resnik 1997), thus maintaining an analogy with the formal sciences.

In many ways my project can be seen as an attempt to provide a more viable realist alternative to the mainstream mentalism or conceptualism of the generative

⁵As pointed out to me by David Pitt, the Platonism of Katz and Postal is by no means the only game in town. I could call the Platonism I plan to argue against KP-Platonism (as in Katz-Postal Platonism) as opposed to, say, a minimally Platonic view (such as Linsky and Zalta's account discussed in 4.3) but since the former is the standard Platonist position in linguistics, I will just refer to KP-Platonism as Platonism for simplicity.

programme without the pitfalls of Platonism. For instance, I will maintain the claim that linguistics concerns sentences as types and not physical tokens. In so doing, I accept that a central question to the discovery of the true foundations of linguistics concerns the appropriate “level of abstraction” of linguistic objects (à la Katz 1996). I will also assume that there is some kernel of truth to the popular infinity claim or rather that the posit of linguistic infinity is not easily escapable for a realist, despite it being potentially irrelevant for the more desirable property of linguistic creativity. Lastly, I will argue that linguistics is indeed distinct from the study of linguistic competence. However, I will not go as far as to identify it with the study of abstract mind-independent objects as *per* Platonism.⁶

In terms of the overall aim of the thesis, this chapter serves as a necessary precisification of the notion of “structure” thus providing an account of the realism aspect of the “structural realism” which was advocated as an interpretation of the previous part.

In the next section, I draw from the Platonistic (and nominalistic) critiques of generative grammar in identifying the essential characteristics of a given realist theory of linguistic foundations. In section 9, I hope to show that Platonism cannot meet this challenge. In section 10, I describe an alternative view which can meet these requirements and in addition offer a more naturalistic account of the foundations of linguistic theory and its objects. This account offers a mathematical structuralist foundation for linguistic theory in which linguistics is a science of natural languages conceived as quasi-concrete structures (in terms of Parsons 1990). In order to tailor the mathematical structuralism of Shapiro (1997) and Resnik (1997) to linguistics, viewed here as a semi-empirical enterprise (or semi-formal, depending on your perspective), I will have to banish the dogma of abstract objects, in the form of the type-token distinction, currently pervasive within the philosophy of language and linguistics. This is the primary task of section 10.4.

⁶However, in section 8.3 I will endeavour to clarify the notion of “mind-independence” to a certain extent.

Chapter 8

Three Desiderata

In this section, I shall outline three important properties of natural language that any realist theory of linguistics should respect. Most of these properties are familiar from various critiques of the generative or biolinguistic programme. The first argues for a central place for the concept of linguistic infinity, despite being potentially irrelevant for linguistic creativity. The second concerns the so-called correct “level of abstraction” for the objects of linguistic theory, namely sentences. Lastly, the final property deals with the relationship between a grammar as theory of linguistic structures and a theory of competence and while it denies their identification (in line with Platonism), it also argues for a particular account of their interaction (in line with Devitt 2006), namely that linguistic competence has to respect aspects of the structure rules of the grammars and *vice versa*.

8.1 Linguistic Creativity and Infinity

Creativity

One of the most discussed properties of natural language is that of linguistic creativity. Despite being assumed to be a universal component of cross-linguistic competence, the notion has not always been clearly described. Part of the problem is that the phenomenon of creativity has not always been separated from the concepts and terms used to model it, such as “linguistic infinity”, “discrete infinity”, “generatively enumerable” etc. This is a mistake and has led to much of the confusion behind the concept of creativity and its place in linguistic theory.

Infinity issues have dominated the foundations of linguistics debate and often informed the rejection or acceptance of various frameworks. For instance, as Searle notes “[w]ithin structuralist assumptions it is not easy to account for the fact that languages have an infinite number of sentences” (1974: 4). For years, Chomskians have placed the need for a computational system with recursive elements at the forefront of their syntactocentrism and the generative programme on the whole. Katz (1996)

argued that due to the infinity of natural language, both Bloomfieldian nominalism and Chomskyan conceptualism fail as interpretations of linguistics because there are simply not enough concrete tokens to capture the generalisations of grammatical theories (essentially restaging the debate between nominalists and Platonists within the philosophy of mathematics). Langendoen and Postal (1984) produced a proof to the effect that the cardinality of natural language exceeds generative capacities, and thus standard accounts of competence, in being of the same magnitude as a proper class (see section Part I section 2.2 for details).

It seems that paradigms rise and fall at the feet of this central linguistic explanandum. However, it is not at all clear what linguistic creativity is or even if it requires linguistic infinity (and in fact the contrary has been convincingly argued by Pullum and Scholz (2010)). I will not rehash this entire debate here, I will however try to make sense of the creativity claim and determine to what extent it goes hand-in-hand with the theoretical posit of infinity. My conclusion will be that infinity should in principle be accommodated within an account of the science of linguistics for reasons other than those usually offered for creativity, but only if one is to be a realist.

What is linguistic creativity? To say that the literature is unclear would be an understatement. However, a natural starting point to this discussion might be found in the comments of Chomsky who placed special significance on this property in the discipline. For instance, consider Chomsky (1964) and (1966) respectively.

The central fact to which any significant linguistic theory must address itself is this: a mature speaker can produce a new sentence of his language on the appropriate occasion, and other speakers can understand it immediately, though it is equally new to them (Chomsky, 1964: 50).

The most striking aspect of linguistic competence is what we may call the 'creativity of language', that is, the speaker's ability to produce new sentences that are immediately understood by other speakers although they bear no physical resemblance to sentences that are 'familiar' (Chomsky, 1966: 74).

There are a few things to notice about the above quotations. The first is that there is no mention of the concept of infinity in either. Given that the expressions which language users actually encounter can only constitute a finite magnitude, the above characterisations potentially allow for an upper bound on the capacity to produce new sentences i.e. a limit to creativity. The second thing to note is the idea that creativity so conceived involves the cognitive ability to interpret novel expressions without prior analogy. Note the emphasis of "new" or unfamiliar sentences here. Part of the reason behind this insistence is to block Hockett-like accounts involving creativity by analogy. Hockett (1968) attempts to cast doubt on the ubiquity of linguistic creativity by suggesting that corpus data indicates that most sentences encountered in daily life are merely variations of a more commonly used/heard set of

sentences (perhaps a precursor to contemporary Construction Grammar accounts?). Chomsky, however, is careful to distinguish the creative “use” of language from the “creative aspect of language” itself. The former may indeed be constrained by various limitations but the latter allows for much more freedom of expression, at least in principle (see Chomsky 1982 for discussion). Nevertheless, freedom of expression still puts us quite significantly shy of infinity claims.

Consider a statement from Chomsky (1972) in which the concept of “indefiniteness” of size surfaces.

Having mastered a language, one is able to understand an indefinite number of expressions that are new to one’s experience, that bear no simple physical resemblance and are in no simple way analogous to the expressions that constitute one’s linguistic experience” (Chomsky, 1972: 100).

Again in the above quotation, empiricist or analogy-based accounts of creativity are explicitly blocked but the idea of an “indefinite number” of expressions is also introduced, which amounts to a denial the possibility of a fixed upper bound on creativity. It is at this stage that one may be tempted to introduce infinity into the picture. However, we are still some distance from requiring linguistic infinity for the notion of creativity under discussion.

Consider the example, presented in Pullum and Scholz (2010), of a standard *haiku*. A haiku typically involves 3 lines with a maximum of 17 syllables (5 in the first and last lines and 7 in the second). The possibilities for haiku creation are clearly finite, yet seemingly “indefinite” in the required sense (somewhere in the region of 10^{34} in Japanese). As Pullum and Scholz note, “the set is large enough that the competitions for haiku composition could proceed continuously throughout the entire future history of the human race [...] without a single repetition coming up accidentally” (2010: 127). This is meant to be a case that shows that infinity is not a necessary condition for creativity. We can see that if we relax the parameters on composition, the cardinality of the creative capacity increases dramatically, yet we are still well-within the bounds of the finite.¹

A similar sentiment on the separation between creativity and infinity is suggested in Evans (1981).

It is unfortunate that Chomsky’s writings have led people to equate the creativity of language use with the unboundedness natural languages display. Linguistic creativity is manifested in the capacity to understand new sentences, and the speaker of a finite language such as the one I have described can manifest it (327).

Evans provides us with a simple language (with 20 axioms linked to a finite vocabulary and a composition axiom) which (similar to a haiku case) allows for a wide

¹Cf. Hockett (1968) for a similar example involving base-ball scores.

range of combinatorial expression (100 sentences) and a disposition towards the understanding of novel expressions. So far, we seem to have a few core components of an account of linguistic creativity, of which infinity is not one. These components include, genuine novelty in terms of non-analogy, indefiniteness in number and flexible composition.² It seems to me that all of these features can be comfortably accommodated by means of the principle of compositionality.

We have seen the principle in previous chapters, but it essentially states that the semantic value of a complex expression is determined by the semantic value of its parts and their syntactic combination.³ Clearly creativity is an important property of natural language and any theory of linguistics, realist or otherwise, should be able to account for it.

So at which point does infinity enter into the picture? The usual story is linked to recursion, iteration and merge. However, I think it goes deeper than these specific mechanisms to the very idea of rule-following in linguistics and the philosophy of language.

Rule-following and Infinity

In this section, I hope to show that realism places an added burden on linguistic theory in terms of infinity claims than do strictly physicalist frameworks. The idea that the theories of natural language are provided by rule-based grammar formalisms has held sway since the seminal *Syntactic Structures* (1957). Two related ideas informed both the inception of formal language theory and the centrality of syntax within the generative tradition in general. The first is that a language can be seen as a collection of sentences of finite length over a finite vocabulary and secondly that a grammar (viewed as a theory of language) generatively enumerates the sentences of that language. Chomsky (1959) goes on to add “[s]ince any language L in which we are likely to be interested is an infinite set, we can investigate the structure of L only through the study of the finite devices (grammars) which are capable of enumerating its sentences” (137). The rules or functions which we specify for a given language are informed by the specific constructions of the natural language under study.

Natural languages such as English notoriously allow for iterative constructions such as those involved in conjunction, subordinate clauses and adverbial modification. Consider the examples from Pullum and Scholz (2010: 114) below.

It is evident that *I exist* is a declarative clause, and so is *I know that I exist*, and so is *I know that I know that I exist*; that *came in and went out* is a verb phrase coordination, and so is *came in, turned round, and went out*, and so is

²Technically, “indefiniteness” is not a property of Evans’ example or the Haiku case.

³The literature on compositionality is much too vast to go beyond a quick statement here. Suffice to say that almost every aspect of its definition is up for grabs. See Shieber and Schabes (1991) for a promising account in terms of synchronous grammars.

came in, saw us, turned round, and went out; that *very nice* is an adjective phrase, and so is *very very nice*, and so is *very very very nice*; and so on for many other examples and types of example.

The idea is that at no non-arbitrary point can we stop the chain of grammatical constructions or rather at no stage in the sentence production can we say ‘this is no longer English’. Thus, natural language seems to be “closed” under recursive rules such as the rules characterising the constructions mentioned above. In this way, we are confronted with a parallel of the Sorites cases in the philosophy of language. Given the nature of certain vague predicates such as *bald* or *tall*, we cannot determine the point at which the predicate disapplies to an object (which can have effects on the validity of rules such as *modus ponens* or principles such as bivalence in certain systems used to model the phenomenon).⁴ If indeed we are dealing with “closure” principles as in first-order logic (FOL), then the generated set (or ‘theory’ in the logical sense) would be unproblematically and denumerably infinite. However, in the case of natural languages, things are generally not this precise. The recursive rules of formal languages do not perfectly capture the nature of natural languages and their constructions. If they did, then there would be no difference between formal and natural languages, but there clearly is such a difference.⁵ Natural languages are sloppy and imprecise, their rules are malleable and violable. More controversially put, there might indeed be a point at which a further iteration of *very* yields an ungrammatical sentence (to borrow a phrase from David Pitt, we might “generate ourselves out of the language”). Nothing I am saying here depends on taking ‘grammatical’ to be a vague predicate (although I think ‘acceptable’ certainly is).⁶ The point is that recursion might indeed be a useful element of the grammars we use to model natural language constructions but it is not a necessary feature of the languages themselves, *mutatis mutandis* for infinity.

An important element of the above characterisation and connection with Sorites series is that of natural languages as concrete objects and linguistic rules as modelling something in the messy physical world. However, if we accept that linguistics is in part a formal science, concerned with some type of abstract objects, similar to mathematics and mathematical logic, this limitation is lifted. On this account, the

⁴I thank Henk Zeevat for suggesting this possible connection to me.

⁵Indeed, some theorists, such as Montague, embrace this claim. But this view is far from generally accepted nowadays.

There is in my opinion no important theoretical difference between natural languages and the artificial languages of logicians; indeed, I consider it possible to comprehend the syntax and semantics of both kinds of language within a single natural and mathematically precise theory (1970: 398).

⁶It is interesting to note that Boolos (2000) entertains this possibility for the axioms of set theory and the existence of large cardinals, when he writes of the subset axiom in ZFC “[b]ut it does not seem to me unreasonable to think that perhaps it is not the case that for every set, there is a set of all its subsets (267)”. Thus, mathematics itself may not immune to these cardinality worries and vagueness might seep into our notions of even the most precise of sciences.

rules of our grammars specify not only model the features of our natural languages, much like the syntactic rules of, say, propositional logic (PL) specify the *wff*'s it generates. If sentences are not constituents of mental states or concrete tokens, then we are free and indeed required to treat the rules of our grammars as determining the structures of our languages. Generativists themselves often make use of this formal analogy, for instance consider Pinker (1994: 86).

By the same logic that shows that there are an infinite number of integers – if you ever think you have the largest integer, just add 1 to it and you will have another – there must be an infinite number of sentences.

As Pullum and Scholz correctly counter, the case for the discrete infinity of the natural numbers is established by the axioms of Peano arithmetic which include a successor function (and an induction axiom schema), and there is no analogy of this operation in the case of natural languages. But a die-hard Platonist (or realist) could insist that there are other mathematical avenues available to arriving at the requisite cardinality (denumerable infinity or \aleph_0). Perhaps one could avail oneself of the idea of weak limit cardinals which do not require anything like a successor function to arrive at denumerable infinity. Postal (2003) has a somewhat nuanced argument for the connection between natural numbers and natural languages. He argues, by *reductio*, that if one assumes an upper bound on an iterative series of sentences in English, then one can show that its logical implications (that the iterations stop at sentence m rather than $m+1$ or $m-1$) cannot be met. The above reasoning is meant to show that the posit of an upperbound on the set of sentences is to be rejected (this is compatible with my suggestion above that such constructions are vague not infinite, if we are to be naturalists).⁷ Nevertheless, the realist has no principled reason for rejecting the idea of closure operations in natural language nor that of languages as sets or collections of expressions (as Chomskyans genuinely adherent to the concept of 'l-languages' are wont to do). The original 'vastness result' of Langendoen and Postal (1984) is testament to the limits of logico-linguistic reasoning. Returning to Katz (1984),

[G]rammars are theories of the structure of sentences, conceived of as abstract objects in the way that Platonists in the philosophy of mathematics conceive of numbers [...] They are entities whose structure we discover by intuition and reason, not by perception and induction (18).

⁷I am not sure that this argument necessarily entails infinity. Following Hockett (1968), consider the rules of any baseball game (which include time constraints). It is easy to see that for any real game, the ultimate score could always have been higher or lower than it in fact was but this does not mean that the score of any baseball game is potentially infinite. I think the analogy here is not with the denumerable infinity of the natural numbers but rather with their 'countability' which can be finite in set theory (i.e. a finite subset of \mathbb{N}). In addition, it assumes that no sequence of sentences has a maximal length.

On this view, natural languages themselves are systems of these sentences and the rules of the grammars governing their interaction are proven in the same way as we would prove theorems in number theory such as Fermat's last theorem.⁸ Thus, linguistic infinity should be an element of any realist account of linguistic ontology and the foundations of the science, notwithstanding its relation or lack thereof to creativity. If recursion is an aspect of our best linguistic theory (grammar) then recursive structures are aspects of linguistic reality. If the set of sentences of a given natural language is closed under conjunction or other recursive operations, then much like the case for formal languages such as PL or FOL, NL is discretely (and trivially) infinite. In section 9.1. I will discuss how a realist might escape a strict reading of this infinity requirement while maintaining the rule-following commitment.

8.2 Of Tokens and Types

Another core component of the realist persuasion in linguistics is the emphasis on the correct "level of abstraction" for the interpretation of its theories. Originally presented in Katz (1984), it has undergone some variation and revision in Katz (1996) and Postal (2003, 2009).⁹ Thus there are a number of related strands to this line of reasoning and I hope to do them justice in this section.

The idea can be summarised as follows. The same species of problem that befell the nominalist or American structuralist project affects the biolinguistic or conceptualist one, namely they were pitched at an insufficient level of abstractness.

Thus, with conceptualism [mentalism], as with nominalism, there is a possibility of conflict between a demand that grammars satisfy an extrinsic, ideologically inspired constraint and the traditional demand that grammars meet intrinsic constraints concerning the successful description and explanation of the grammatical structure (Katz, 1984: 195)

In order to correctly meet the "intrinsic" constraints such as infinity, recursion and structural hierarchy, the psychological level (or "extrinsic ideology of mentalism") is inadequately abstract on this view (another way to understand the quotation above is that grammars do not need to go beyond 'descriptive adequacy'). Therefore, we need to ascend to a higher level of abstraction to capture these linguistic properties. In the absence of a systematic correspondence between the formal structure and the physical system, an extreme interpretation of this problem could be expressed

⁸For example, proving that $a^n + b^n = c^n$ is true for any positive integers where $n > 2$ might be a similar task to proving $a^n b^n c^n d^n$ where $n \geq 1$ is a string not accepted by a context-free grammar. But the former is certainly a different task from showing that Swiss-German is not such a language (see Shieber 1985 for details).

⁹The question of the correct level of abstraction has also received an innovative treatment in Bromberger (1989). More on that in section 10.4.

as a charge of a category mistake at the heart of the biolinguistic movement (or “incoherence” in Postal (2009)). Thus, a physical system (a human brain) is not something capable of possessing properties such as infinity (or capable of veridical description in terms of the set-theoretic merge operation). Mental states and physical tokens cannot be recursive or infinite, only sets and other mathematical objects are amenable to such description.

The problem does not disappear with the limitation of structure either. In the minimalist program (1995), Chomsky investigates the *minimal* structural requirements needed to explain the gulf between the child’s initial state and the adult’s later competence, as well as language evolution. This marks a departure from the often complex linguistic architectures of the Extended Standard theory (*circa* 1970) and Government and Binding (1981) which posit various levels of representation and interfaces between these levels. Once again, the central desiderata of linguistics is to account for the perceived discrete infinity of linguistic expression and the hierarchical nature of syntactic organisation. According to minimalism, in order to explain these features, one need only posit a binary merge function which takes two syntactic objects and outputs one. Technically, there are two merge operations, *external merge* which takes two distinct objects as input and *internal merge* which allows embedding and thus allows for recursion. Furthermore, *internal merge* involves duplicating items within the operation. For instance, if we merge syntactic objects α and β to form the unordered set $\{\alpha, \beta\}$ and there is a γ such that γ is a member of α and we merge this object with $\{\alpha, \beta\}$, we would have two copies of γ in the resulting structure. In this way, we are supposed to be able to account for all the usual movement operations with very minimal apparatus in the syntax (and various constraints on the operations).

Merge, however, is set-theoretic in nature. The universe of set theory (nonde- numerably captured by the universe V) generally takes sets to be outside of space and time, finite or infinite and abstract. Before continuing, it is important to clear up one potential confusion here. The objection is not supposed to be that mathematical models are being used to describe a physical system. This is a commonplace practice in science and does not presuppose that all mathematical modelling generates incoherent ontologies (as I have shown in the previous chapters). The reason for the specific problem in the biolinguistics tradition can be couched in terms of the lack of a systematic correspondence between elements of the model and elements of the target system, i.e. we have no idea how elements of the set-theoretic operation of merge correspond to neurophysical structures. On the one hand, we want to explain discrete infinity, recursion and syntactic hierarchy through the all encompassing set-theoretic operation of merge. On the other hand, we want to provide a naturalistic explanation of language in terms of the human brain and biology. Postal (2009) believes that these requirements pull in opposite directions and thus cannot be met in the same object simultaneously, namely an I-language. Thus biolinguistics is stuck with an untoward or “incoherent” ontology (at least at its current stage). Or as Postal (2003: 242) puts it “[t]he received view claims that an NL is something

psychological/biological [...] a state of an organ [...] And yet it has been unvaryingly claimed in the same tradition at issue that NL is somehow infinite. These two views are not consistent".¹⁰

The move made by Platonists then is simply to raise the level of abstraction of sentences to that of sets and other abstract objects, thereby proffering a coherent ontology for the interpretation of linguistics. Returning to Katz (1984),

Sentences, on this view, are not taken to be located here or there in physical space like sound waves or deposits of ink, and they are not taken to occur either at one time or another or in one subjectivity or another in the manner of mental events and states. Rather, sentences are taken to be abstract and objective" (18).

Postal (2009) presents a similar argument to this effect. However, he follows Katz (1996) in availing himself of the type-token distinction. If linguistic theory or grammars were indeed about brain-states etc. as the biolinguist would have it, then the sentences of these theories would have to be at the level of tokens, not types (which are here conceived of as abstract objects). There are two issues with this position, he claims. For one thing, it seems out of touch with linguistic practice in which grammars usually deal with "island constraints, conditions on parasitic gaps, binding issues, negatively polarity etc." (Postal, 2009: 107). Importantly, these accounts are rarely, if ever, informed by evidence from neuroscience or psychology (as one would expect if they were truly concerned with brain-states). Therefore, he concludes that these accounts are concerned with sentence types conceived abstractly.¹¹

Sentence tokens exist in time and space, have causes (e.g. vocal movements), can cause things (e.g. ear strain, etc.). Tokens have physical properties, are composed of ink on paper, sounds in the air [...] Sentences have none of these properties. Where is the French sentence *a signifie quoi?* – is it in France, the French Consulate in New York, President Sarkozy's brain? When did it begin, when will it end? [...] Such questions are nonsensical because they advance the false presupposition that sentences are physical objects (Postal, 2009: 107).

¹⁰I have shown this reasoning to be deeply flawed in Part I.

¹¹For Katz (1996) the abstractness concern in linguistics is a special case of the general problem of abstractness in the formal sciences. An account such as the strict finitism or "inscriptionalist nominalism" characterised by the Hilbert programme, for instance, failed as an appropriate interpretation of mathematics according to Katz. In order to capture the infinity of mathematics via the empiricist scruples of nominalism, only reconstructed language about the infinite is permitted, "mathematics is about mathematical expressions" (Katz, 1996: 273). The objection is simply that to make sense of such talk, we need either expression types, which take us back to abstract objects, or expression tokens, which need to allow for unactualised *possibilia* which in turn are no less metaphysically suspect than abstract objects. Katz, however, neglected the vast literature on actualist reinterpretations of quantified modal logic, some varieties of which posit *contingently nonconcrete* objects in an attempt to avoid commitment to *possibilia*.

These considerations lead Platonists to conclude that linguistics is concerned with sentences on the level of abstract objects, in the sense of non-spatio-temporally extended entities. Truth in linguistic theory or in its grammars is then determined by correspondences between the sentences of the theory and these objects. I think this latter leap is not required and in fact rather inimical in light of better options and tamer ontologies, both of which I shall present in sections 10 and 10.4. However, suffice to say that there is some kernel of truth to the notion that linguistic grammars and the theories they inform do possess a formal and abstract level of description through the analysis of sentence types (or whichever type of basic unit with which one begins). Furthermore, a realist account of linguistics should provide an appropriate interpretation of this aforementioned level of abstraction and linguistic practice as it is. We will return to this issue in some detail in section 10.4.

8.3 Realism and Respect

So far, I have not said much about what I take “realism” to be exactly. This was a somewhat intentional move on my part. The properties or desiderata of the previous sections stemmed mostly from Platonist critiques of mentalism. Thus, they pushed a specific agenda and ontological attitude. The next series of arguments stem from a very different ontological approach to linguistics, similar in its focus on *concreta* to conceptualism but in line with Platonism in its rejection of representationalism or the idea that speakers of a language represent/know/cognise the grammar rules of their language. The chief proponent of what is called the “linguistic conception” (as opposed to the “psychological” of generative grammar) is Michael Devitt in his book *Ignorance of Language* (2006). Devitt is a realist but certainly not a Platonist. Hence, by appreciating his stance and its intersection with those of Katz and Postal, I think, we may be able to carve out the realist position in more detail.

Given what I have said above, we might be tempted to consider realism to be non-ontologically committing. And in so far as we would be tempted to do so, I think we would be correct. Linguistic realism, as I conceived of it, is simply the position that linguistics is *about* something outside of psychological reality. Theories or models of language, i.e. grammars, tend to describe this extra-mental reality and not the linguistic competence of speakers. It is in saying something more precise about what exactly this non-psychological reality is that realists diverge. Platonists hold that it is an abstract extra-physical reality, while nominalists, such as Devitt, prefer a physicalist account (my own account will draw from aspects of both ontologies). Another way of putting this point is that both Platonistic and nominalistic realists hold that language *qua* object of linguistic inquiry is not in the brain but where they say it is differs quite drastically from one view to the next.

At this juncture, some distinctions between different *mind-independence* claims might be helpful (I follow George (1996) here, although the thrust of his distinctions

is more epistemic than ontic). Let us consider three options and their interrelations.

MI_1 The subject matter of the field is constituted by “entities distinct from minds” (George, 1996: 297).

MI_2 There are truths of certain elements of the domain that are (in principle) unknowable.

MI_3 Elements of domain do not depend for their existence on minds.

Disciplines which exemplify the first kind of mind-independence are general physical theories such as cosmology, astronomy, biology, chemistry etc. MI_1 is compatible with the epistemic claim of MI_2 . This is how George cashes out Chomsky’s claim that settling all the facts about the mind will settle all the facts about natural language, namely by denying MI_1 while allowing for the possibility that such complete knowledge is inaccessible to us. MI_3 is a distinct claim from MI_1 . For instance, Platonists affirm MI_1 , MI_2 and MI_3 while Devitt and nominalists of a similar persuasion only affirm MI_1 . It seems then that realism only requires adherence to something like MI_1 .

So given the above characterisation of realism, unsurprisingly, part of the next desideratum of a realist account of linguistic foundations will be the rejection of competencism or the view that linguistics concerns the psychological states of language users (what Devitt calls “the psychological view”) or what I have called MI_1 above. I will follow Devitt one step further in adding another aspect to this desideratum and that is a notion of RESPECT between the posits of the grammars and the processing rules of competence.

[A] theory of a competence must posit processing rules that respect the structure rules of the outputs. Similarly, a theory of the outputs must posit structure rules that are respected by the competence and its processing rules. (Devitt, 2006: 23).

This is what Devitt calls the “Respect Constraint”. A few things before we consider it more carefully and why it belongs within a realist conception of language. One of Devitt’s favourite examples is that of von Frisch’s theory of the “waggle dances” of bees. He uses the theory to make three general distinctions (in the spirit of realism). Von Frisch observed that bees use a form of communication called a “waggle dance” to indicate the direction and distance of food sources to other bees in the hive. For example, if a bee returns from a food source over 100 metres away, it will employ a “waggle dance” (a “round dance” if less than 100). The angle at which the bee arrives in the hive reflects the angle with relation to the sun of the bee’s path from the food source while distance is indicated by the speed of the dance.

Devitt uses this example (and others) to distinguish between (1) the theory of the waggle dance (a snapshot of which I provided above), i.e. the behavioural outputs of

the bee, and the theory of the bees' competence in its execution. Von Fisch's theory clearly only provides insight into the former. Another distinction is between (2) the structure rules of the dance, which can be diagrammatically presented easily, and the structure of the processing rules of the individual bees themselves (i.e. what's going on when they compute various distal and directional parameters for communication) of which we have no conception. Last is (3), the respect constraint or the claim that "the bee's state of competence, and the embodied processing rules that constitute it, must "respect" the structure rules of the dance in that they are apt to produce dances that are governed by those rules" (Devitt, 2008: 205). It seems clear that von Frisch's theory of bee dances, grammar of their language if you will, is concerned with the structures of the dance itself as per (1) and not the structures of their competence or performance of it ((2)) of which we know nothing except that it respects the rules of the theory in the sense of (3).

From the above distinctions, Devitt claims that grammars of linguistics are true of linguistic reality and not human psychology (where English, French or isiZulu are our waggle dances). From this conception of grammars he defines his minimal position (M) below.

A competence in a language, and the processing rules that govern its exercise, respect the structure rules of the language: the processing rules of language comprehension take sentences of the language as inputs; the processing rules of language production yield sentences of the language as outputs (Devitt, 2006: 57).

The onus is on the generativists or conceptualists to prove that we need more than this minimal posit, i.e. prove that representationalism is correct. This has been a notoriously difficult task, in most cases representationalism was merely assumed (compare Fodor's (1981) glib term for the mentalism of Chomsky as simply 'the Right View'). In addition, early psycholinguistics was initially meant to determine the connection between the processing rules of performance and the grammar rules of competence. This was generally considered to be an unsuccessful venture (even by its own proponents at the time). Nevertheless, it is not my concern here to challenge Devitt's position from a conceptualist or mentalist perspective (see Collins 2007, 2008a, 2008b; Lawrence 2003; Rey 2006b, Slezak 2007 for such arguments).

The central question with relation to this research is 'if we are realists, why not stop with Devitt's linguistic conception?' Unfortunately, there are some problems with the view in light of the other desiderata I consider and general issues about which a realist might be concerned. For one thing, in this chapter (and Devitt's book) a lot is said about what linguistics is *not* about but so far we have not delved into the question of what linguistics *is* about and here lie the problems for the realist position I take in this part of the thesis.

In the preface to *Ignorance of Language*, Devitt describes both his initial fascination with and initial resistance to linguistics. He states (of his thoughts during his

graduate years) “[s]urely, I thought, the grammar is describing the syntactic properties of (idealized) linguistic expressions, certain sounds in the air, inscriptions on paper, and the like [...] It rather looked to me as if linguists were conflating a theory of language with a theory of linguistic competence” (Devitt, 2006: v). This thought is apparently the seed out of which the main ideas of the book grew. Now most realists would agree on the last statement, in fact Katz (1981), (1984) and Postal (2003) stress the fallacy of conflating the knowledge of language and language itself allegedly present in generative linguistics. It is the first claim, that grammars are about “sounds in the air” and “inscriptions on paper”, that seems to be at odds with realism. Once again, we seem to be at the wrong level of abstractness. Concrete tokens are insufficiently abstract for the interpretation of most of what linguists do. We saw in the previous section that there is some kernel of truth to the type talk of Platonists and in so far as “idealised token” means type, we are fine but I doubt that this is what Devitt has in mind.¹² To reiterate, grammars, on this view, describe structure rules which constitute representational systems outside of internal mental representational systems (but are respected by them). As we saw with the above characterisation of (M), sentences are supposed to be inputs for processing and sentences are also outputs of processing, but what are sentences? Are they physical tokens, “inscriptions on paper” or “sounds in the air”, surely not since this would not sufficiently interpret the practices of actual linguists as *per* Postal’s objection in 8.2.

At this point, I think realist and conceptualist objections converge to a certain extent. Ludlow (2009) claims that “while Devitt purports to be offering a proposal that is faithful to linguistic practice, the range of linguistic phenomena and explanation he surveys is limited” (394). This limitation cannot, for instance, deal with postulates of covert material in syntax (which have no phonological expression), such as PRO (also see Collins 2007, 2008a). If our structure rules concern physical tokens (sounds, writings etc.) then elements which do not overtly appear through these media pose a problem. Much of linguistic practice and methodology involves the use of assumed entities or items (See Parts I.3 and II.4). Katz (1971) linked the Chomskyan revolution in linguistics to the Democritean revolution in early scientific thought in that it aimed to expose the underlying reality behind appearances.

Let us consider another example from contemporary (generative) linguistics. In the literature on negative concord (NC), where the meaning of a negated expression involves a balance of negative elements, covert material tends to show up quite frequently in the analysis. Compare the following sentences, one from English (a double-negation language) and the other from Spanish (a negative concord language).

¹²Devitt (2008) writes “according to my “linguistic conception” a grammar explains the nature of linguistic expressions. These expressions are concrete entities external to the mind, exemplified by the very words on this page” (249). In section 10.4, I will offer some alternatives which might capture this intuition more adequately.

(8.1) I didn't not go to work today.

DN: I went to work today.¹³

(8.2) María no puede encontrar a nadie

María *not* can find to *nobody*

NC: Maria can't find anyone.

In order to account for NC in a way that offers a unified analysis of negation, Zeijlstra (2004) starts with the claim that “NC is analyzed as an instance of syntactic agreement between one or more negative elements that are formally, but not semantically negative and a single, potentially unrealised, semantically negative operator” (Biberauer and Zeijlstra, 2012: 345). More specifically, Zeijlstra defines negative concord as a type of AGREE relation between a formally semantically interpretable [iNeg] feature and at least one uninterpretable [uNeg] feature. Thus, NC languages can contain elements which only look negative but actually bear the [uNeg] feature. In other words, some negative elements on the surface can be semantically non-negative in reality. In addition, this AGREE relation is a Multiple Agree relation which means that multiple [uNeg] elements can be c-commanded by one element bearing [iNeg] in the feature checking. Finally, it is argued that in grammatically justified situations, a covert [iNeg] can be assumed to c-command any overt [uNeg] and “of course, language-specific properties determine whether this non-realisation possibility is actually employed” (Biberauer and Zeijlstra, 2012: 349). Therefore, the NC agreement is between one formally and semantically negative operator (which is often covert) and one or more overt non-semantically negative elements. Now look at an example from Czech in which it is argued that no overt negative elements are at play in the negation (of the surface syntax).

(8.3) Dnes nikdo *nevolá*

NC: Today nobody calls

[DnesOp_{-[iNEG]}[TPnikdo_[uNEG]nevolá_[uNEG]]]

In (8.3) nothing in the surface form of the sound and written tokens in Czech produces the negation by itself (according to this analysis at least).¹⁴ The grammar then assumes a covert operator to generate the negative meaning. Thus, linguistic grammars have to be about something other than the physical tokens and their

¹³English speakers do make use of a form of understatement called “litotes” which also involves double negation but not always for the sake of retrieving a strong positive reading as in the example above. Litotes is largely pragmatic.

¹⁴This analysis is supported by the impossibility of double negation in Czech (and similar languages) and the cross-linguistic typology of possible negative configurations put forward by Zeijlstra and others.

structural relationships in order to account for this and similar research. Remember, Devitt is not aiming to provide a revisionist conception of the foundations of linguistics and therefore his account needs to square with contemporary practice.¹⁵

On a related note, in section 8.1, I argued that realists have to take posits of the grammars (and their consequences) to be actual features of linguistic reality. Such posits include recursive structure rules and closure principles which seem to lead to infinity claims. Thus, either we need to be able to ascend beyond the level of physical tokens which fail to interpret such claims or provide a naturalistic account of infinity claims in linguistics (I attempt to do both below).

Lastly, despite the issues with the nominalism of this proposal, we will incorporate an element (or two related elements) from its realist core, namely that linguistics is the study of language not the study of linguistic competence (or knowledge) directly and that the study of language and the study of competence needs to be connected by a RESPECT constraint (the latter is the specific contribution of Devitt's account). After all, we do produce and understand natural languages and it would be strange if we could not account for this aspect of the human experience in linguistics, even if it is the independent study of language systems conceived of in a realist manner. In other words, the study of natural language does have an empirical element which needs to be addressed. More on this dilemma in the next section.

8.3.1 Taking Stock

So far, I have been attempting to determine the key aspects of a realist account of linguistics. I have argued that although potentially unrelated to creativity (which requires compositionality), linguistic infinity cannot be ignored by realists (as it can potentially be jettisoned by linguists of other persuasions). I affirmed the need to ascend beyond a level of physical tokens or mental ones in the interpretation of grammatical theory. Lastly, I accepted that linguistics is the study of a competence-independent (mind-independent₁) linguistic reality but I restricted this claim by insisting (with Devitt 2006) that this reality be linked to linguistic competence *via* a structural respect constraint. For clarity, I provide the list below as a guide for the ensuing discussion.

1. A realist interpretation of linguistics ought to (a) account for creativity in terms

¹⁵Devitt's (2008b) response to these worries is to dismiss them as "highly theoretical" and "abstract" objects of syntactic theory distinct from the convention fixing (communicative) intentions of speakers. In addition, he claims that PRO and the like are not likely to be determined by innate cognitive constraints or UG (he vacillates somewhat on this position later). I have two issues with this response. One is that, as realists, we should care about theoretical posits of grammars (see section 8.1) and the practices of linguists as *they are* if we are truly to be interpretationalist (and not revisionist). Then secondly, the issue over whether such covert structure is determined by innate elements of UG is beside the point for a realist. If it is in the grammar then we should be able to explain it.

of novelty, compositionality etc. and (b) account for the potential infinity (denumerable or otherwise) of natural language(s).

2. Linguistic theory is a theory of sentences at the level of types or more generally realism needs to pitch linguistic theory at the correct level of abstraction.
3. (a) Linguistics is the study of natural language, not the study of the knowledge of or competence in that language, and (b) grammatical structures (and rules) need to be respected by the structures of competence and *vice versa*.

What remains to be shown is that Platonism is not the best way of capturing these three conditions on a realist account of linguistic foundations. This is the topic of the rest of the paper. In the next section, I will show that Platonism fails on counts (1) ((a) and surprisingly (b)) and (3b). Its failure on count (2) will have to wait until section 10.4 for explicit treatment.

Chapter 9

Against Platonism

In this section, I will be rather brief since my argument is straightforward (and additional arguments against Platonism can be found in Part I, section 3.3 and in section 10.4). Simply put, Platonism is not the best way of capturing the three desiderata or conditions on a realist account of linguistics as described above. I will start with an argument to the effect that Platonism cannot account for either creativity or the kind of infinity usually associated with linguistics. Then I will argue on the basis of Benacerraf's famous dilemma for mathematical truth (1973) that the respect constraint cannot be met by Platonists in any plausible way and therefore as with mathematical Platonism a gulf is created between the truth of our linguistic theories and our knowledge of this truth (competence). Lastly, I will make a general claim (following Soames (1984)) that mathematics (as well as logic) and linguistics are conceptually distinct and if indeed linguistics is a formal science, it is a *sui generis* one.

Before we get to this task, however, let us review what the Platonist position is. Essentially, Platonism is an account of linguistic foundations which holds that linguistics is the study of abstract mind-independent objects. The Platonist takes all of the syntactic and semantic structure posited by grammars not merely as useful tools for describing mental states or physical tokens (i.e. models of linguistic reality) but as constituting an independently existing linguistic reality directly. A natural language, like a formal language, is an abstract object in the sense of being non-spatio-temporally extended and comprised of sets of sentences. On the view we have been considering (that of Katz and Postal), sentences are ontologically similar to numbers, sets and geometric figures.¹ Natural languages are simply systems of these sentences, describable by us through reason and intuition *a priori*.

¹Postal (2003) states that "an NL is a set-theoretic object, a collection, in fact, a bit more precisely, a collection of sets, where each set is a complex object composed of syntactic, semantic, and expression objects. The traditional term for these sets is "sentence," so that it is appropriate to say that an NL is a collection of sentences" (237).

9.1 The Right Kind of 'Wrong View'

An important aspect of the linguistic Platonist position (dubbed the 'Wrong View' by Fodor (1981)) is that it contends that there is a static universe of natural languages (and the sentences of which they are comprised) already existing independently of human beings and language users. We discover languages, we do not create them. Much like numbers and sets exist independently of mathematicians who study them or the bean counters who use them, if there were no speakers or users of natural languages, there would still be natural languages and sentences.

Once this metaphysical point is appreciated, I think Platonism's incompatibility with the type of creativity discussed previously can be gleaned. In section 8.1. we looked at creativity in language and its role in linguistic theory. I argued that it involved the use and appreciation of novel sentences (to the user), the manipulation of composition rules and the indefiniteness of the number of expressions for which it allowed. The problem is that according to Platonism every sentence of every language already existed (or exists in an atemporal sense) before they were used or thought of. The mere instantiation of existing objects through production or comprehension is surely not to be considered novelty? A child counting to a previously uncounted number might be performing an impressive feat but it would not be deemed 'creative' in the sense that the term is used in linguistics. The number existed prior to the child's recitation and the child was merely its mouthpiece.

We should pause to appreciate the depth of this Platonist idea. Every sentence of *Ulysses* or the *Odyssey* of Homer (or every other book in every language which has ever existed) existed in a very definite way before Joyce or Homer ever set pen to paper (or voice to word). Perhaps they were the first to pluck these particular sentences from the heavens but this activity can hardly be called creativity. And if we are redefining creativity in light of this view, then we should at least admit that the subject has changed from the concept discussed by Chomsky, Evans and others. Certainly, Platonism can accommodate an impoverished notion of novelty-to-a-speaker similar to the new number-to-a-counter but the stronger notion (involving genuine creation) would be inaccessible on this ontological account.

Platonists might want to bite the bullet on this one. But I think that it is related to a more pervasive misinterpretation of linguistic methodology as is evinced by the wholly unexpectedly problem of linguistic infinity. Part of the motivation for linguistic Platonism was to better capture infinity claims and the 'vastness of natural language'. In fact, Katz (1996) argues that without Platonism, the vastness result of Langendoen and Postal (the Cantorian proof that the cardinality of natural language is the size of the continuum or a proper class discussed in section 2.2) does not go through as an objection to generativism or competencism. It seems as though Platonism and infinity go hand-in-hand conceptually. So how then, can I claim that Platonism is at odds with linguistic infinity?

My contention is that the infinity with which linguistic Platonism provides us is

the wrong kind of infinity for linguistics, which is usually underpinned with a rather more constructivist approach to infinite expression. Before I present my case, it is important to remember that we are in the interpretation game not the revision one. It is in the spirit of that aim that I argue that linguistic infinity is not to be understood *statically*, as *per* Platonism, but rather *dynamically*, as *per* constructivism (or even strict finitism).

A brief history of the foundations of mathematics might be in order here. Constructivism, or intuitionism, starts with the idea that mathematics is the product of human thought and therefore should be accessible to human mental capabilities. Iemhoff (2015) describes Brouwer's initial conception as follows.

The truth of a mathematical statement can only be conceived via a mental construction that proves it to be true, and the communication between mathematicians only serves as a means to create the same mental process in different minds.

A famous example of intuitionist thinking is the rejection of rule of double-negation in classical logic which states the following equivalence: $\neg\neg p \equiv p$. Intuitionistic logic rejects this rule because the proof of the negation of a negated sentence is not the same as a positive proof of the sentence or as Heyting put it "a proof of the impossibility of the impossibility of a property is not in every case a proof of the property itself" (1956: 17). One consequence of the above reasoning is the failure of the law of excluded middle in intuitionistic logic. The reasoning goes that since there are statements in mathematics (such as the Continuum hypothesis or the Riemann hypothesis) for which there is neither a positive proof nor a refutation (nor a clear path to either), and since having a refutation means being able to show the positive proof false, the principle cannot hold in every case. The underlying intuitionistic move responsible for the various departures from classical logic mentioned above (and beyond) is the link between truth and knowability present in the framework.²

The notion of proof and construction appear within this redefinition of mathematics through the relocation of the human mathematician to the subject role in the mathematical process. For example, in Hilbert (1899) the claim "one can draw" in geometry is taken to be synonymous with "there exists". Here again we see why classical principles such as excluded middle fail. Existence claims in intuitionism are equivalent to the production of exemplars and there are certain claims (such as the Continuum hypothesis etc.) for which we cannot do so (nor produce refutations).³ This is in turn coupled with a mentalistic approach to construction. As Heyting notes,

²George (1996) describes this connection in terms of MI_2 above or "Brouwer, and the constructivists generally, do affirm the mind-dependence₂ of mathematics" (297). That is, settling all the facts about the mind might indeed answer all of the questions of mathematics.

³I am indebted to comments and emendations provided by Bernhard Weiss for this section.

Isolating an object, focusing our attention on it, is a fundamental function of our mind. No thinking is possible without it. In isolating objects the mind is active. Our perception at a given moment is not given as a collection of entities; it is a whole in which we isolate entities by a more or less conscious mental act (1974: 80).

Naturally, much of the philosophical motivation behind constructivism and intuitionism centered around the concept of infinity. The idea of an infinite series incapable of comprehension in its entirety by a human mind was contrary to the core precepts of this position. For instance, instead of starting with the successor function and the axioms of Peano arithmetic, for the intuitionist the natural numbers start with the process of counting. According to Heyting, this is the mental process of isolating perceptions of entities and then creating more of these entities in one's mind (and in time, importantly). A fuller survey of intuitionism in mathematics is unfortunately outside the scope of the present work. I do, however, want to draw a comparison between this picture of mathematics and the initial idealisations of the nature of linguistics as a science. Shapiro offers us a helpful way of thinking about constructions.

I propose that we think of the constructions as performed by an imaginary, idealized constructor, obtained in thought by extending the abilities of actual human constructors. Then we can sharpen dynamic language and the various "construction problems" by articulating exactly what abilities are attributed to the ideal constructor (1997: 184).

The idea is that we can interpret dynamic talk of "constructing" mathematical objects, or following mathematical rules, in terms of these ideal mathematicians not limited in the same way as actual mathematicians are. Thus, certain moves might still not be permitted by intuitionists (such as inferring p from $\neg\neg p$) but we are also not stuck in the very literal readings of such talk (bound by actual performance). Compare this to the much quoted opening lines of Chomsky's *Aspects of a Theory of Syntax*.

Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its (the speech community's) 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 this language in actual performance. (Chomsky, 1965: 4).

Indeed, much of the talk surrounding the concept of generative grammars, recursively enumerable sets and discrete infinity is constructivist in linguistics. An ideal speaker is capable of expressing an infinite number of sentences of her language (has a generative grammar in her mind), but the infinity in question is a constructive

not an actual one.⁴ It is the product of mental competence, it is a mental activity like counting is for intuitionists such as Brouwer and Heyting.⁵ The ideal speaker is following a procedure set out by the rules of her grammar or “the language, in that sense, provides instructions to performance systems” (Chomsky, 2000). In addition, with this understanding of infinity, novelty can also be rescued. We, as human language users, genuinely create the structures of our languages as we produce and comprehend them. Thus, new sentences can be produced by following certain rules (the rules of the grammar of our language). The sentences do not exist prior to these constructions. In this way, linguistic infinity is understood as an infinite capacity to produce sentences of the language (as it should be) rather than mysterious access to an (atemporally) existing infinity of objects as Platonism seems to suggest.

9.2 Benacerraf’s Dilemma and Respect

The failure of the respect constraint, I argue, is due to another larger issue with Platonism in the philosophy of mathematics. The problem was famously identified by Benacerraf (1973) and has significantly altered the landscape in the foundations and philosophy of mathematics since. The dilemma posed by Benacerraf makes the claim that the quest for mathematical truth pulls in two opposing directions with relation to a uniform semantics and a(n) (causal) epistemology. The argument takes the form of placing two demands on any theory of our knowledge of mathematics. Namely, that

- (1) the concern for having a homogeneous semantical theory in which semantics for the propositions of mathematics parallel the semantics for the rest of the language, and (2) the concern that the account of mathematical truth mesh with a reasonable epistemology (Benacerraf, 1973: 661).

Benacerraf held that all (or most) accounts of mathematical truth fail to find the appropriate balance between these two demands, in fact more than that, the demands seem inversely proportional in these accounts. Consider Platonism. In providing a standard truth-conditional semantic account which dovetails with the semantics for the rest of language, Platonists avail themselves of reference to abstract objects. In other words, the truth of mathematical statements about numbers, sets and the like is determined by their correspondence to abstract entities, non-spatio-temporally extended, in a similar way to how reference to physical objects is supposed to be

⁴In the spirit of the previous parts and grade one of mathematical involvement, one might say that a constructivist model of mathematics does a better job of reflecting mathematical reasoning than does a Platonist one.

⁵Pylyshyn (1973) makes similar comparisons between Chomsky and intuitionists like Heyting. Chomsky himself states that “[o]ne could perhaps take the intuitionist view of mathematics as being not unlike the linguistic view of grammar” (1982: 16).

fixed (in a Tarski-style semantics). However, in providing such a semantic account, we cannot begin to make sense of our causal contact with the former objects (by definition) and thus are left with no (causal) account of our mathematical knowledge the likes of which we have for ordinary physical objects. In the opposite direction, empiricist accounts of mathematical knowledge tend to root it in the familiar physical causal world (the Hilbert programme or Devitt's analysis for linguistics) but fail to then specify how the necessary truth of these objects is obtained in a uniform semantics for ordinary discourse.

This is not the place to go into too many details about Benacerraf's dilemma, but suffice to say that by endorsing Platonism for linguistics, Katz and Postal essentially accept its lot.⁶ Postal (2003) admits that "[a] formal, abstract object-based view of linguistic ontology, of course, faces the classic epistemological problem often raised in connection with mathematics and logic of how knowledge of abstract objects can be obtained" (251). He defers discussion, however, to Katz' *Realistic Rationalism* (1998). We will get to a discussion of some of these ideas below but for the sake of this dialectic I would like to recast Benacerraf's dilemma in terms of the "respect constraint" discussed in the previous section.

In the previous section, in accordance with Devitt (2006), I advocated for the need for a realist condition on the relationship between the structure rules of grammars and the structures of linguistic competence (whatever these may be). This move was made in part to "ground" realist accounts of linguistic theories (of the outputs of language comprehension and production) in the mental activities of language users and *vice versa*. In relation to this point, I further argued for an interpretation of all talk of infinity and generative grammars in terms of constructivist mathematics. In a sense, this condition was suggested (imposed) to prevent language from getting away from us, if you will.

The issue with Platonism in linguistics is that, much like the Benacerraf problems for Platonism in mathematics, its ontology pulls in an opposite direction *vis-à-vis* the respect constraint. More precisely, if the structure rules of the grammars designate objects in a Platonic realm, i.e. abstract objects without spatial or temporal dimensions, then how are we to account for their relationship with the physical competence of language users in their use or acquisition of such objects? In other words, how does the abstract ontology of linguistic Platonism account for our knowledge of language, i.e. our linguistic competence? Furthermore, if we take linguistic constructivism seriously, there might indeed be mathematical structures which are incapable of being comprehended by a human mind but surely there are no such

⁶Katz's (1995) response to this dilemma utilises what I call an *argument from linguistics*, to dismantle Benacerraf's case. He argues that surface form is not always a guide to deep structure (by means of the famous *eager to please* versus *easy to please* case) and that causal theories of knowledge are not the only game in town. His idea is that Platonism needs neither a uniform semantics for countenancing its objects nor a causal theory of knowledge. Unfortunately, in the absence of concrete proposals on either side, this position is hard to evaluate.

linguistic structures (so called impossible grammars are not part of human language by definition). We cannot impose the condition that competence respects the structures of linguistic reality if it is possible that this reality completely outstrips human comprehension. In the other direction, why would linguistic grammar rules or the structures they posit *qua* abstract objects need to correspond in any way to real world constraints any more than higher-order set-theoretic entities should respect our abilities to conceive of them? In this way, intuitionism in mathematics can be interpreted as the attempt to establish a RESPECT constraint on mathematical theory and the mental competence from which it is spawned. In addition, Lewis (1975) can be viewed as attempt to establish a similar constraint between grammars and conventions of linguistic communities.⁷ Nevertheless, whatever the status of Platonism is for mathematics, it poses a particular problem for understanding or respecting the relationship between natural languages and the speakers (or knowers) of these languages. From the rather tame realist separation of linguistic reality from linguistic competence or knowledge of language, Platonism effectively creates a gulf between them.

9.3 Conceptual Distinctness

In the previous subsections, I aimed to show that Platonism cannot meet my first and last desiderata of a realist theory of linguistic foundations. In this section, I will briefly concern myself with another corollary of the Platonist view of linguistic objects. This is the view that given realism, linguistics itself must be a formal science on par with mathematics and logic. In order to show this reasoning to be fallacious, I will apply a similar (realist) strategy employed by Soames (1984) to the effect that linguistics is not cognitive psychology (i.e. my first desideratum).

The strategy proceeds in the following way. In order to establish that two types of theories are *conceptually distinct*, one has “to show that they are concerned with different domains, make different claims, and are established by different means” (Soames, 1984: 155).⁸ Challenge accepted.

I think the first two requirements are relatively uncontroversial (although potentially question-begging against Platonists), namely that linguistics and mathematics are concerned with different domains and make different claims. Linguists are concerned with natural languages such as English, Swahili and Tamil. They care about the structures of these languages, their cross-linguistic similarities and differences

⁷In fact, respect and the idea of “well-modelling” are not entirely distinct in my view.

⁸Soames also uses the tool of what he calls *empirical divergence*, i.e. linguistic structures are unlikely to be isomorphic to psychological structures, which on the face of it seems to be in contrast to my RESPECT constraint. Empirical divergence, however, is a much stronger claim on the relationship or lack thereof between linguistic theory and the theory of competence, and respect certainly does not require anything as strong as a morphism or structure-mapping to hold.

and how they change over time. When linguists write grammars for specific languages or attempt to model certain formal properties of various constructions, they are constantly required to make sure that their grammars and properties correspond to actual languages spoken (or signed) in the world. This is accomplished sometimes by means of checking linguistic intuitions (their own and those of other native speakers) or corpus data. To put the point somewhat differently, the linguistics practiced on a planet of speakers cognitively and socially distinct from humans, might look very different from our own, or at least the grammars and constructions might (the linguist's job might still be the same though). Linguists might ask the same questions but the content of their answers would be different. On a standard Platonistic account of mathematical theory, this is not the case. Set theory on earth looks exactly the same as set theory on Pluto or Mars (even if they were populated with different sorts of creatures).⁹ I think that this is generally the case because the two types of theories are "established by different means". Mathematicians consult their intuitions *a priori* while linguists are bound by certain contingent linguistic phenomena and behaviour (at least in part).

Linguists, like empirical scientists, might use mathematics (as in formal language theory and truth-conditional semantics) as tools or even essential tools but this is different from mathematics as method. Even in its strongest form, the disanalogy persists. Without sets, functions, morphisms etc. linguists might not be able to describe linguistic reality (or competence). But there is a difference between saying "we can't describe-without-mathematics linguistic reality" and "we can't describe linguistic reality-without-mathematics". This is the Berkeley fallacy, mentioned by Yablo (2013), that statements like "we can't imagine a tree non-perceptually" do not entail statements of the form "we can't imagine an unperceived tree" (1016). I would opine that linguistic research constitutes, at most, the use of (perhaps essential) mathematical tools but not necessarily mathematical methodology.

Furthermore, mathematical methods are different from tools. The methods of mathematics involve things like postulation, induction, implicit definition, impredicative definition and construction.¹⁰ Such methods are generally absent from linguistic theorising and grammar construction. The linguist's job is not done after postulating a mathematical possibility, the possibility only becomes linguistic if it is instantiated by some real world language. For decades, research into finite-state grammars was abandoned due to Chomsky's claim to have shown that such formalisms did not concern natural language constructions in any significant way (see Pullum (2011) for the falsity of that claim). For instance, following optimality theoretic phonology,

⁹Of course, these creatures could have a different logic and this might affect the mathematical structures they discover or postulate. But certain structural relations seem to be ubiquitous. Consider group theory which deals with a basic notion of symmetries. By studying the symmetries of structures, we shed light on the nature of these structures themselves whatever they may be.

¹⁰See chapter 5 of Shapiro (1997) for an overview of the place of these methods in the history of mathematics.

we see that certain sequences of syllables are not realised by any human languages such as (C)VC (see Hammond 1997 for discussion). Once this is deemed the case, phonologists are no longer interested in such patterns, i.e. they are not *linguistically* interesting. The task of a mathematician has no such empirical restriction. In opposition to this, Postal (2003) claims that there are natural languages for which no knowledge exists. To be a natural language is just to obey certain constitutive laws and if we can specify an object that obeys these laws and is unlearnable, then there are unlearnable natural languages. This is an implicit definition and a corollary of the Vastness theorem. Still, it is not clear to me why learnability is not one of the constitutive laws of natural languages as formalisability might be for their formal counterparts. “Learnability” here, I mean “first-language acquirable” since it is quite possible to learn non-natural languages (perhaps formal languages or programmes).¹¹ In addition, allowing for such unrestricted uses of implicit definition violates the *respect* constraint.

Linguistics certainly seems to use mathematical tools in identifying the properties of its objects (as do many sciences) but it does not seem to mathematically define the objects of its inquiry *a priori* or rather use mathematical methods. In Lewisian terms, linguistic and mathematical objects seem to be *orthogonal* to one another or as I have put it (following Soames), the fields are “conceptually distinct”. Of course, one could argue that not all formal sciences are alike and linguistics is unique (a similar line is taken in Katz (1981)). In the rest of the chapter, I aim to lend some credence to this idea.

¹¹On the distinction between acquirable and learnable, or first language acquisition versus second-language, recent studies focused on a linguistic savant named “Christopher”, who otherwise impaired mentally has the miraculous ability to learn numerous languages with remarkable fluency, by Smith and Tsimpli (1995) and Smith, Tsimpli, Morgan and Boll (2010) is particularly interesting.

Chapter 10

Ante Rem Realism

So far I have argued that Platonism (and nominalism) failed to capture certain conditions or desiderata of a realist interpretation of linguistic theory. I proffered these desiderata in accordance with arguments presented for these very positions. What remains to be shown is that there is a realist alternative to Platonism that can account for (1) linguistic creativity and infinity, (2) the appropriate level of abstraction present in current linguistic accounts or grammars and (3) both the separation of linguistic reality from competence and the mutual respect constraint between them.

In the following sections I will describe a view of the foundations of linguistics in terms of a non-eliminative structuralism similar to that offered for mathematics by Shapiro (1997) and independently by Resnik (1997), I call this view *ante rem* realism. I hope to show that the ontology that this position brings with it is coherent in the spirit in which Platonism was offered but does not suffer from the same problems as described in the previous section such as Benacerraf worries. Furthermore, this account allows for a more naturalistic interpretation of linguistics as an empirical science with formal aspects by debunking the various misconceptions associated with abstract objects and the corresponding type-token distinction.

10.1 Mathematical Structuralism

The motivation behind mathematical structuralism can be traced back to Benacerraf and the dilemma he presented (see section 9.2). The core idea of this foundational picture in mathematics is that mathematics is a theory of structures and systems of these structures. In this way there is a shift from the traditional (perhaps) Fregean concept that numbers, sets and other mathematical entities are abstract objects, unencumbered by spatial or temporal properties. The core insight is that it is structures and not objects which are the vehicles of mathematical truth (and knowledge). This presents an entirely different conception of the nature of the enterprise as well as the concept of a mathematical object itself. Structuralism is a broad framework with

historical antecedents ranging from the Bourbaki group and Dedekind to Hilbert and even Benacerraf himself. Thus, there are a number of varieties of the idea at work within the contemporary philosophy of mathematics. I will try to stay as broad as possible for the moment, although I do plan to endorse and develop a particular variety of what is referred to as *ante rem* or non-eliminative structuralism for linguistics in the next section.

In order to understand this view on the foundations of mathematics, we need to answer a few preliminary questions. Firstly, what are structures on this view? And how do they relate to traditional objects of mathematics? Secondly, whatever they are, how do we come to know about them? Then finally how does understanding mathematics as a theory of these structures get us out of Benacerraf-types worries? I hope to provide some potential answers to these questions in this section.

Shapiro starts his book with the slogan “mathematics is the science of structure”. He continues by way of example,

The subject matter of arithmetic is the *natural-number structure*, the pattern common to any system of objects that has a distinguished initial object and a successor relation that satisfies the induction principle. Roughly speaking, the essence of a natural number is the relations it has with other natural numbers (1997: 5).

This holds true for groups, topoi, euclidean spaces and whichever mathematical structure is studied by mathematicians. Let us focus on the natural-number structure for a moment and consider its objects. What is a number on this view? Essentially, it is nothing more than a place in a natural-number structure. The only way to talk about the number 2 or 5 or 4892001 is with relation to other places in that structure, i.e. 2 is the successor of the successor of 0 or the number 2 is the third place (if we start from 0 as Frege did) of a natural-number structure, it is in the second place of an even-number structure and the first place of a prime number structure and so on. The same holds for other mathematical objects, the idea being that these objects are only interpretable in accordance with some background theory. As Parsons puts it, “the idea behind the structuralist view of mathematical objects is that such objects have no more of a ‘nature’ than is given by the basic relations of a structure to which they belong” (2004: 57).

The concept of a group is often taken as a canonical example of a structure. A group G consists of a finite or infinite domain of objects and a two-place function called the group operation. This function satisfies four properties (or axioms). It is associative (associative property), there is some identity element (identity property), it is closed (closure property), and every element in the domain must have a reciprocal or inverse (inverse property). Now there are many different types of groups which mathematicians may wish to study. We could look at finite groups (groups with finite domains) or Abelian groups (groups whose elements are also commutative). The basic group structure is the same and the structure is given to us by the relations its objects

have to one another (according to the four properties). The objects themselves are of no importance to us, they might as well be point-particles, martians, jelly-beans or rice-crispies, it doesn't matter.¹ What matters is the structural relations one object (whatever it is) has to another in the group, we only care about the structures. In fact, we can even talk about structures in isolation from any objects. Shapiro characterises his own position in the following way.

The first [*ante rem* structuralism] takes structures, and their places, to exist independently of whether there are any systems of objects that exemplify them. The natural-number structure, the real-number structure, the set-theoretic hierarchy, and so forth, all exist whether or not there are systems of objects structured that way (1997: 9).

The other versions of structuralism offer similar accounts. They differ, however, in important respects. For instance, the question of whether or not structures can themselves be considered mathematical objects. For set-theoretic structuralists, inspired by model theory, the answer is yes. Structures are set-theoretic entities themselves. For modal structuralists, structures are not objects of study. Hellman (1989) utilises this framework to avoid reference to individual mathematical objects all together (by replacing such talk with talk of *possible* mathematical objects or number-systems in his case), it is thoroughly eliminative. The point is that there is no one answer to the question of the nature of structures themselves, different structuralists will provide radically different accounts. Another question concerns the background logic, which varies from first-order with identity to second-order and modal logic given different accounts of structuralism.

We have looked at the question of what structures are and what traditional mathematical objects are within them, i.e. merely places-in-structures devoid of individual meaning or importance. The last question to confront in this section is how this framework aims to avoid Benacerraf's dilemma. Recall that Benacerraf's claim was that the more uniform the semantics, i.e. the more the objects of mathematics were treated on par with the objects of ordinary discourse, the further we get from a tractable epistemology. The semantic problem was that we were forced to the treatment of abstract objects as singular terms referring to non-spatio-temporal entities. This created an ontological gap untraversable by standard causal accounts of knowledge. But with structuralism, there is no such reference since there is an ontological difference between an object and a place in a structure. Neither numbers nor sets commit us to individual abstract objects (as with Platonism),² but merely to places-as-objects in

¹Compare this to the description of a category in category theory. "A category is *anything* satisfying the axioms. The objects need not have 'elements', nor need the morphisms be 'functions' [...] we do not really care what non-categorical properties the objects and morphisms of a given category may have (Awodey, 1996: 213).

²Although they are referential in a manner consonant with ordinary discourse as I will show in the next section.

natural-number structures or set-theoretic structures. The existence of these kinds of objects is provided by the axioms (as we saw with group theory) or relational properties of the structures. These axioms and structural relations, in turn, can be known by us in a presumably more sound epistemic manner.³ I shall leave matters here for now and more details will follow when we consider a specific structuralist proposal for linguistics in the next section. But before getting to that, I think a brief detour into the metaphysics of structures is in order.

Ontological Dependence

So far I have described the beginnings of the structuralist interpretation of mathematics. I suggested that some of the answers to ontological questions within its remit are specific to individual structuralist frameworks. I think we can do a bit better than this. Therefore, in this section, I will discuss one particular ontological claim that serves to provide further insight into an essential element of various structuralist approaches and perhaps even separate talk of structures from talk of ordinary objects. In section 10.2, I will argue that this core aspect of the structuralist programme in mathematics, namely the specific notion of dependence (as described by Linnebo 2008), can be found in the linguistic project in both syntax and semantics as well.

Structuralism is often characterised in contrast with Platonism, especially in terms of its treatment of objects. For a Platonist, abstract objects are analogous to ordinary physical ones in that they are ontologically independently of one another. My toaster no more relies on my backpack than my carpet relies on my desk chair for its existence. Mathematical objects, on the other hand, have no such independent existence according to structuralists. These objects *qua* positions in structures depend on other positions for their very existence and on the structures as a whole. For this reason, Linnebo (2008) distinguishes between two notions of dependence.

ODO Each object in D [domain of some mathematical structure] depends on every other object in D (67).

ODS Each mathematical object depends on the structure to which it belongs (68).

The difference between ODO and ODS is that the former just says that the objects in a structure depend on other objects such as some natural numbers depending

³Of course, knowledge of axioms also results in further epistemological questions but of a much different order to knowledge of Platonic objects. For example, for Gödel, the truth of axioms of set theory “force themselves upon us” so much so that “despite their remoteness from sense experience, we do have something like a perception of the objects of set theory” (Boolos, 2000: 266). Boolos attenuates this extreme claim somewhat to suggest that perhaps only certain axioms have the desired effect (e.g. extensionality and pairing). Parsons (1980) attempts to pick up on the “perception” analogy for mathematical intuition and claims that there is indeed a phenomenon which answers to it.

on other natural numbers while the latter adds that the existence of one object in a structure ensures that the structure itself exists or is ensued by the existence of the structure as a whole. Another way to think of ODS is that structures are ontologically prior to positions (which explains Shapiro's quote above stating the possibility of uninstantiated structures). Linnebo goes on to argue that requirements such as non-circularity (cashed out in terms of well-foundedness) militate against ODO straightforwardly and (perhaps) ODS to a lesser extent. Notwithstanding various difficulties with either or both of these dependence relations, it is important for my purposes that (ante rem) structuralism incorporates a strict notion of *upwards dependence*.⁴ Upwards dependence is the relation in which objects depend on the overarching structures as opposed to depending on their own constituents.⁵

The ubiquitous and dominant definition (or family of definitions) of compositionality in philosophy of language and linguistics relies on both ODO and ODS. It usually takes the form of something like the following.

The meaning of a complex expression is determined by the meaning of its constituents and their method of combination.

In this picture, the meaning of linguistic objects is a function of the meanings of their constituents and the overarching syntactic structure in which the expression finds itself.⁶ Although, the principle of compositionality is usually not understood in

⁴As Linnebo suggests, this might be a fundamental difference between the realm of the mathematical and the realm of the physical. In keeping with this distinction, the view of linguistic reality advocated here endorses a hybrid ontology of linguistic structures.

⁵Of course, ODO might be deemed necessary in cases in which the dependence on the entire structure might lead to contradiction such as ordinal set theory, in which the dependence on the totality of sets is notoriously problematic. I thank Stephen Read for pointing this out to me.

⁶Contemporary inferentialism challenges this claim. Instead of an atomistic view of compositionality as the one cited above, it proposes a holistic view essentially based on the concept of implicit definition. Atomism (and standard compositionality) presupposes that individual constituents have meaning independently and these meanings combine to yield the meanings of the complex expressions in which they are contained. This is directly analogous to the mathematics case. In fact, standard compositionality is based on the compositionality of formal languages such as propositional and predicate logic (which have straightforward homomorphisms between the the algebras constituted by the rules of the syntax and semantics respectively). However, the compositionality of inferentialism is different. In a particularly illuminating passage, Peregrin describes an important aspect of inferential roles.

Thus roles are given merely through an 'implicit definition', and just as Quine (1969, p. 45) claims that 'there is no saying absolutely what the numbers are, there is only arithmetic', we can claim that *there is no saying absolutely what inferential roles are, there are only rules of inference (and compositionality)* (2015: 53).

Here we see that, just as in the mathematical structuralist case, the inferential roles which determine the meanings of sentences and words are *upwards dependent* on the linguistic structures (or social-normative networks) in which they are found. Thus, ODS holds. It seems that the top-down or sentential approach of which Brandom (2007) speaks is equivalent to the upwards dependence of *ante rem* structuralism.

terms of existence of meanings but rather meaning or semantic value individuation. In the next section, I will argue that despite certain departures from pure structuralism syntactic and semantic structures generally rely on something akin to ODS.

10.2 Linguistic Structures

Linguistic Structures and Realism

Previously, I described a general framework, neither obviously Platonist nor nominalist in nature, which confronted Benacerraf's dilemma by eliminating the need for reference to ontologically occult abstract objects. Importantly for our purposes, the *ante rem* structuralism of Shapiro and Resnik is a realist theory of the foundations of mathematics. As Shapiro states, "as articulated here structuralism is a variety of realism" (1997: 6). He distinguishes between two kinds of realism within a model-theoretic semantics (such as Tarski's). 'Realism in ontology' or the idea that singular terms in the language of mathematics denote mathematical objects which genuinely exist and 'realism in truth value' which states that grammatical sentences in mathematics have definite truth values (either true or false). He claims that his version of structuralism is realist in both senses.

In order to appreciate the realism of this proposal, one has to delve into the notion of an "object" –as a position in a structure– which it incorporates. The claim is that natural language provides as with two uses of the concept. In the one more frequent case, we treat positions as offices or roles, which are multiply realisable in terms of entities. For instance, some uses of *President* or *rook* are examples of these cases. They do not denote individual objects as in *The President has the right to overrule the senate* or *The rook can move three places*. Shapiro calls this 'places-as-offices'. There is another sense of the term in which we treat positions not as the offices or roles they occupy but as genuine singular terms denoting objects. Examples are sentences such as *The President had lunch with the Dalai Lama today* or *The rook ate the queen at d7*. This is the 'places-are-objects' perspective. *Ante rem* structuralism takes this latter concept as primary. Of course, as Shapiro notes "[w]hat is an office from one perspective is an object –and a potential officeholder– from another" (1997: 11).

Now from the above, we can see how this form of structuralism is realist in ontology and realist in truth value. In arithmetic or number theory we take numbers to be objects, but in set theory they are offices. Consider the number 2, "[i]n one system, [finite von Neumann ordinals] $\{\emptyset, \{\emptyset\}\}$ occupies the 2 place, and in the other [Zermelo numerals] $\{\{\emptyset\}\}$ occupies that place" (Shapiro, 1997: 11). In either case, the numeral 2 is a name picking out an object *qua* position in a structure and statements involving the numeral are true or false but in neither case are we committed to an individually existing number in the Platonic sense. All we need is for the structure

to exist (and there are various ways of ensuring this, see chapter 3 of Shapiro (1997) and section 10.1 below). In fact, this example presents one of the advantages of this theory over Platonism. According to Platonists, numbers are individual mathematical objects and mathematical objects are sets. If this is the case, then there is a fact of the matter as to which sets constitute the natural numbers. But von Neumann ordinals and Zermelo numerals have different set-theoretic consequences for numbers, since on the former account ' $2 \in 4$ ' is true while on the latter it is not. How do we decide which theory is correct? With structuralism we don't have to decide, since both theories are true in virtue of being concerned with the same natural-number structure, not the individual numbers and their correspondence to specific abstract entities or individual sets.

10.3 Quasi-Concreteness

The account I offer essentially makes use of the same claim to realism as in the mathematics case. If "mathematics is the science of structures", then linguistics is the science of linguistic structures. *Ante rem* realism is the position on the ontology of language that states that linguistics is concerned with abstract patterns or structures and grammars are theories or rather models of those structures. My account does, however, depart from that of Shapiro (and Resnik) in significant ways. Consider the following remark made by Resnik concerning linguistics.

"Take the case of linguistics. Let us imagine that by using the abstractive process [...] a grammarian arrives at a complex structure which he calls *English*. Now suppose that it later turns out that the English corpus fails in significant ways to instantiate this pattern, so that many of the claims which our linguist made concerning his structure will be falsified. Derisively, linguists rename the structure *Tenglish*. Nonetheless, much of our linguist's knowledge about *Tenglish qua* pattern stands; for he has managed to describe *some* pattern and to discuss some of its properties. Similarly, I claim that we know much about Euclidean space despite its failure to be instantiated physically (1982: 101)"

In linguistics we seem to be concerned with a specific class of structures, those which are instantiated in the real world. These are the structures that are produced by human linguistic competence, i.e. the outputs of competence. In this way, I amend the structuralism of Shapiro to include what Parsons (1990) calls quasi-concrete objects. These objects or positions-in-structures, in my view, are comprised of a mixed ontology. Parsons offers the existence of such objects as an objection to pure structuralism but I see no serious reason for why this cannot be compatible with it for the case of linguistics (Shapiro himself takes this concept as a friendly amendment). Parsons states that there are "certain abstract objects that I call quasi-concrete, because they are directly 'represented' or 'instantiated' in the concrete" and he includes

as an example of such an object “symbols whose tokens are physical utterances or inscriptions” (1990: 304). The idea is that there is an additional relation to the axioms of certain structures that goes beyond pure structuralism, a ‘representational’ (or instantiation) relation. The problem is that these sorts of objects require a representation (or instantiation) relation which cannot be accommodated in the purely structural picture involving nothing additional to intra-structural relations. Or rather

What makes an object quasi-concrete is that it is of a kind which goes with an intrinsic, concrete “representation,” such that different objects of the kind in question are distinguishable by having different representations (Parsons, 1990: 34).

I will return to this point in section 10.4.3. For now, suffice to say that I think that this third kind of ontological category merely marks the boundary (which is vague) between the structures of pure mathematics and those of applied sciences in which I place linguistics.

In *Realistic Rationalism* (1998), Katz offers a similar account for what he calls “composite objects”. Examples of objects like the equator or impure sets (which have physical objects as members) push him towards accepting a third metaphysical category of objects. These are not just objects with dualist parts or feet in both worlds but they stand in a “creative” relationship with one another, i.e. their composition creates a new object distinct from either part.⁷ For instance, the equator is neither a perfect circle nor a line that exactly bisects the circumference of the earth, since “[i]t didn’t exist before the earth was formed and will cease to exist when the earth ceases to exist” (Kaufman, 2002: 219). In terms of impure sets, in *Skeptical Linguistic Essays* (2003), Postal identifies classes of sentences, involving direct discourse, whose ancestral elements actually include physical objects.

This entails that the sets that comprise NL sentences must be able to contain as members or submembers something that can instantiate the endlessly distinct physical properties involved in direct speech. The only way I see that this can be the case is if direct speech segments involve sets that contain the physical properties themselves and not, as in the case of more standard (regimented) linguistic elements, symbols that represent instructions (to a fixed physical apparatus) to produce physical things (Postal, 2003: 193).

8

My account in some ways corresponds to the position Katz and Postal suggest at times despite differing significantly from the one they endorse. Furthermore, I

⁷This creation relation vastly overgenerates and thus in the end fails to maintain the concept of a concrete object since concrete objects stand in indefinitely many relations to abstract objects. See Kaufman (2002) for details.

⁸Impure sets so defined are redolent of ur-elements in set theory.

think that this is a very intuitive picture of the science of linguistics. The idea of quasi-concrete objects is also not necessarily metaphysically occult. We seem to encounter these objects on a daily basis. Consider Boolos' comments to that effect below.

Numbers do not twinkle. We do not engage in physical interactions with them, in which energy is transmitted, or whatever. But we twentieth-century city dwellers deal with abstract objects *all the time*. We note with horror our *bank balances*. We listen to *radio programs: All Things Considered* [...] Some of us write *pieces of software* [...] And we draw *triangles* in the sand or on the board. Moreover bank balances, reviews, palindromes, and triangles are "given" to us "in experience," whatever it may mean to say that (2000: 265).

What Boolos calls "abstract objects", I call *quasi-concrete*. And "what it means to say that" they are "given to us in experience" is just to say they have either instantiation or representation relations in the concrete.⁹ The difference between Boolos' list and linguistic (and some mathematical) objects is that many of the abstract objects on his list are fully determined by the physical objects to which they relate whereas linguistic objects, as I conceive of them, have a generally structural nature in addition to concrete instantiation or representation.

Linguistic Structures and Dependence

What after all is syntax, if not the study of the structural relationships between sentences and their subphrases? Of course, these structures should be additionally exemplified by real world languages but this is merely the addition of the respect constraint for which I argued earlier. The syntax of a particular language is an abstract object much like the *University of St Andrews*. Following Ryle, we cannot ask where the university is exactly since it is the organisation of different ever-changing units, it is a quasi-concrete structure. The positions various buildings occupy could change, the chemistry building could house the biology faculty at some stage and thus change its assignment, some buildings can be removed and others erected. If the entire structure is destroyed, then it no longer exists *in toto*. But it existed once in a temporal and partially physical sense. The syntax (and semantics) of a particular natural language is similarly abstract, it is the organisation of linguistic units or sentences in terms of their structural relationships to one another. If the language dies, so do the systems (physical instantiations of structures) which governed it. Of course through records we could still study the language on a more abstract/formal level as with the University blueprints, we could even resurrect the

⁹Within the context of linguistics, Stainton seems to describe a similar class of objects. "There is another sense of *abstract*, however –namely, things that are not inside the mind yet are not concrete particulars either. They are neither fish nor fowl. Let me coin the term *abstractish* for these" (Stainton, 2014: 6). Within this list he mentions objects very similar to those found in Boolos' catalogue above.

language based on the structures as in the case of Hebrew. Hale (1987) assumes that natural languages, like mental states, have temporal parts notwithstanding their lack of physical dimensions.

In terms of section 10.1, the above reasoning culminates in the claim of upwards dependence for linguistic structures (or adherence to ODS). Individual nodes in a syntactic tree are dependent on the structural configuration of the entire tree and defined in terms of it. Take for instance, the adjunct-complement distinction in syntax. What makes a constituent an adjunct versus a complement is determined by the position it takes in the structure. Adjuncts are usually iterable and positioned external to the head or main clause while complements are limited in number and follow the head directly (or are included in the phrasal structure of the head). Furthermore, empty categories, as discussed in Part II.4., are defined purely in terms of the overall structure of the expression or constituent.

However, ODS cannot capture the fact that intrinsic features of lexical categories can affect the syntactic and semantic structures in which they are present. In other words, linguistic objects cannot be defined in purely structural terms. Hence, the move toward quasi-concrete structures. An example of this phenomenon in semantics is selectional criteria for lexical items. In syntax, subcategorisation picks out this phenomenon. The idea is that some words or phrases require specific structures or arguments for their completion by features of their internal “nature”. Predicates exemplify selectional restrictions and certain verbs require certain kinds of categories for completion. Consider the examples below.

(10.1) # Alude is rusting.¹⁰

(10.2) *Thabo works the car.

In (10.1), the predicate “is rusting” cannot take human agents but usually takes only (metallic) inanimate objects as a subject. Otherwise, it is syntactically well-formed (i.e. the structure is fine). In (10.2), the verb *work* requires a prepositional phrase (optionally, since the sentence is felicitous without a complement too). Thus, the nature of the specific word places requirements on the syntactic structure itself *contra* pure ODS.¹¹ In order to succour pure structuralism, it might be tempting to exploit ODO to account for selectional restrictions and the like. However, ODO is generally interpreted *globally* and it thus a much stronger relation requiring inter-connection between all the objects of a relevant domain.

Quasi-concrete objects project features of their internal makeup onto the overall structures in which they partake. Thus, their ontology cannot be captured in terms

¹⁰The hash-tag is meant to convey semantic anomaly. In many cases, such as literary contexts, these sentences can receive interpretation.

¹¹So-called defective verbs, such as *beware* which only exists in the imperative, are also interesting violations of ODS.

of purely external relations (such as axioms or implicit definition). The way in which this works for linguistics will be the topic later sections.

Taking Stock Again

We can now see that this account can meet all of the desiderata of a realist theory of linguistics. Linguistic creativity and infinity are easily represented as there are no size limits to the linguistic constructions we employ. In addition, we can avail ourselves of the *dynamic* discourse of constructivists, as the linguistic structures which we create as language users could be conceived of as direct products of our mental faculties, despite being amenable to study independent of those faculties. Much like the natural-number structure could have been created or constructed by initial counting procedures of human agents through abstraction (see Shapiro (1997) chapter 4 for a suggestion and Resnik (1982) for a more speculative account), natural language patterns or structures could have been created by the dual need for thought and communication among human cognisers. The rules of either activity leads to a potential or constructive infinity.¹² In terms of the appropriate “level of abstraction”, we have an arguably more sound account than Platonism offered us. After all, *ante rem* structuralism drew inspiration from the classical position on universals and particulars (as Hellman calls structures on this view “*sui generis* universals”). Unlike the previous dualist picture, we have a potentially naturalistic picture available to us. Linguistic grammars are concerned with sentences as positions-in-linguistic-structures. Immediately, we do not run into Benacerraf-type worries about how we as physical beings use abstract objects like sentences if they are not extended in spacetime. Sentences, like numbers, have purely relational and structural components, c-command, governance, scoping relations etc. But unlike numbers, I argue, sentences are part of quasi-concrete structures which include representation relations. In the same sense as the non-eliminative or *ante rem* structuralism discussed above, sentences on this account are *bona fide* objects (in the places-as-objects sense) and linguistic statements concerning them have definite truth-values. Thus, sentences are not to be taken as tokens or “words on a page” and “sounds in the air” or mental states for that matter but abstract objects conceived as places or positions in linguistic structures which are in turn represented or instantiated by those tokens.

Once again, the emerging picture seems rather intuitive in light of actual linguistic practice. Consider a determiner phrase (DP). On most syntactic accounts, it is a structurally designated linguistic item in a hierarchical structure or tree and any

¹²The research of Simon Kirby is especially interesting with relation to this point. Kirby (1999) designed a series of experiments to computationally test the emergence of structure in a population over time with the result that “[t]he simulation results [...] show that compositional, recursive language emerges in a population which initially has no language [...] Purely through the process of being repeatedly mapped from an internal form as a grammar to an external form as utterances and back again, language evolves” (Kirby, 1999: 14).

word or object (sometimes nothing as in the case of null determiners) can satisfy the position. And whatever is in that position is a DP. In addition, the much-discussed postulation of covert material is usually supported by structural reasoning in linguistics, i.e. something must be there since this structure requires it or it stands in a structural relation to something else (recall the negative concord example above). The UG hypothesis itself can be considered structuralist in that it aims to discover the underlying structures of the human faculty of language, the particular items or objects of various languages are rendered inconsequential (this is often a criticism of the claim). Furthermore, consider Jackendoff's Parallel Architecture (discussed in Part II.5), a highly modularised account of the language faculty which consists of various individual generative systems with interface principles or relations between them. On this view, the syntax is not the only generative system (as it is with traditional generative accounts) but semantics and phonology are systems (or "a collection of objects with certain relations" (Shapiro, 1997: 73)) in their own right. The interfaces are concerned with the structures, i.e. the systems at a higher level of abstraction, where non-relational elements are ignored. In addition, the model-theoretic perspective in both syntax and semantics is explicit about the structural treatment of natural language. As Pullum states in the case of syntax,

The grammatical expressions of a human language such as English, together with their syntactic structure (and recall that I regard them as actually existing objects with inherent structure), can be idealized mathematically as relational structures [...] We could ask: What is the simplest and most elegant set of axioms that is satisfied by those structures that are appropriate representations for grammatical English sentences, and thus in effect characterizes grammatical well-formedness for English? (2013: 497).

With relation to realism, one significant advantage of this foundational framework is that it can provide an answer to Quine's (1972) famous challenge to Chomsky concerning equivalent grammar formalisms. Quine's challenge was initially posed to a conceptualist framework, i.e. if two grammar formalisms are weakly equivalent (generate the same set of sentences) then how can we divine which one is cognitively realised in the human mind/brain? Similarly for the Platonist, if sentences are sets and two weakly equivalent grammar formalisms pick out the same sets of sentences (sets of sets), how can we tell which sets constitute the language in question? This is essentially a parallel of the arithmetic case involving the finite von Neumann ordinals and the Zermelo numerals (and also Benacerraf's (1965) objection to Quine's version of Platonism). The answer for the *ante rem* realist is analogous, they both pick out the same natural language structure and thus we have no reason to decide between them.

Another related aspect in favour of this view over its Platonist alternative is the level at which languages themselves are pitched. As previously mentioned, sentences are abstract objects for Platonists. But so too are languages as they are defined as

'systems of sentences'. As Carr put it, "while it is perfectly reasonable to assume that sentences are linguistic objects and thus susceptible to such Platonic interpretation, it is rather novel to argue that particular languages [...] should be taken to be objects of linguistic theory" (1990: 123). Generally, the boundaries between external languages like Dutch, English and German are not sharply defined. Likewise, the Platonic claim is that there is a fact of the matter as to which distinct abstract objects (or sets) Serbian and Croatian correspond to respectively. However, languages in this sense are often politically defined and classified (hence Chomsky's initial reservations about E-languages). In general, these types of languages are within the realm of sociolinguistics and not objects of grammatical theory. On the *ante rem* realist account, Serbian and Croatian, Urdu and Hindi and other such cases have structural overlap. The systems of sentences to which our grammars of these languages correspond are the same or similar natural language structures, they need not be identical to achieve this end nor need there be a fact of the matter as to which structures they correspond to exactly.

For the last desideratum, one way in which to satisfy this condition is to treat the representation or instantiation relation of our quasi-concrete linguistic structures as the RESPECT constraint itself. I will present an argument which is compatible with this proposal in section 10.4.3. Thus, one way in which our linguistic structures or patterns could be represented in the physical world is by respecting the rules of our competence and by those same rules respecting the rules of the structures in turn. This could be achieved by persisting with the idea that the quasi-concrete linguistic structures are comprised of sentences which are the output of our linguistic competence but distinct from that competence, like the waggle dances of Devitt's bees. I think that on this view we have even more options than these available to us for capturing the interdependence of structure and mind while maintaining their distinct natures. Furthermore, if linguistic structures are the outputs of competence and competence is within the evolutionary order of things in the physical world, then given the RESPECT constraint, our linguistic structures are also related to a naturalistic story of language evolution.

10.4 Banishing a Dogma

In section 8.3.1, I promised that I would show that Platonism failed to capture the proper type-token distinction and thus failed to place linguistics at the correct level of abstraction. In the previous section, I suggested a less impoverished notion of this distinction in terms of quasi-concrete structures in which "the relation of linguistic types to their tokens (and in general of quasi-concrete objects to their concrete 'representations') is not an external relation" (Parsons, 1990: 337). I claimed that this account of the requisite abstraction level was more in line with the *ante rem* realism I proposed for the foundations of linguistics as well as some comments and

accounts suggested by Katz and Postal themselves. Despite the fact that a mixed ontological attitude towards abstraction is well-supported in the literature (Hale, Parsons, Stainton etc.), a hard-line Platonist could insist that there is no independent justification for jettisoning the clearer traditional account of types as abstract objects and tokens as their physical instantiations. The claim that quasi-concrete structures seem to “go better” with the ontology I propose is not independent reason for accepting these structures nor it is sufficient justification for my earlier claim that Platonism fails to do abstraction justice. In this final section, I will make the case for abandoning the traditional view of types as non-spatio-temporal abstract objects outside of the causal order. First, however, consider these passages cited in both Katz (1996) and Postal (2009).

There will ordinarily be about twenty *thes* on a page, and of course they count as twenty words. In another sense of the word *word*, however, there is but one *the* in the English language; [...] it is impossible that this word should lie visibly on a page or be heard in any voice (Peirce, 1958: 423)

ES IST DER GEIST DER SICH DEN KÖRPER BAUT: [S]uch is the nine word inscription on a Harvard museum. The count is nine because we count *der* both times; we are counting concrete physical objects, nine in a row. When on the other hand statistics are compiled regarding students’ vocabularies, a firm line is drawn at repetitions; no cheating. Such are two contrasting senses in which we use the word *word*. A word in the second sense is not a physical object, not a dribble of ink or an incision in granite, but an abstract object. In the second sense of the word *word* it is not two words *der* that turn up in the inscription, but one word *der* that gets inscribed twice. Words in the first sense have come to be called tokens; words in the second sense are called types (Quine, 1987: 216–217).

Characterisations of objects such as those presented in the quotations above aim to establish a distinction between abstract and ordinary objects. Once this distinction is in place, there are two options for describing the relationship between these respective types of objects. We could go the traditional Platonist route of removing abstract objects from the causal order by stripping them of physical and temporal parts. This is inimical for the reasons we saw in section 9.2. and a host of others. Bromberger’s account in section 10.4.1 shares features with this approach. Another option is adopting a position called ‘Naturalised Platonism’ (Linsky and Zalta (1995)). This position makes the empiricist claim that properties and sets and other *abstracta* are well-within the causal order and knowable *a posteriori*. Kaplan’s view in section 10.4.2 is consonant with this option. In some ways, Quine too falls within this camp by constraining abstract objects through the same principles (such as Ockam’s razor) that constrain other theoretical entities. Still we are left in some confusion as to how we come to know these entities in the first place.

In order to offer a genuinely naturalised account of Platonistic underpinnings and abstract objects, Linsky and Zalta (1995) propose what they call 'Platonised Naturalism'. The aspect of the project that has particular significance for the current discussion is their identification of the genesis of the issues with the Platonistic positions mentioned in the previous paragraph.

We believe that there are two mistakes in that conception: (i) the model of abstract objects as physical objects, and (ii) the piecemeal approach to theorizing about abstract objects (Linsky and Zalta, 1995: 9).

The first prong of this analysis is particularly important here and I think the main issue with the erroneous accounts of the type/token distinction as presented by Quine and repeated by Katz (1996) and Postal (2009) above. Most Platonists (and many other philosophers) take abstract objects to be analogous to physical objects. If physical objects are 'sparse' or discoverable piecemeal, then so are the abstract objects to which they correspond, if physical objects are 'complete' as in have more properties than we know and are entirely physically determinate, then abstract objects are knowable in their entirety and determinate in detail (either true or false for all properties) and lastly if physical objects have 'backsides' or underlying hidden structures, then abstract objects are similarly complex. In some sense, this picture is natural since abstract objects are often determined by 'abstracting' from physical objects. But this dichotomy brings with it serious problems. Linsky and Zalta go as far as to assert that it is the root of Platonism's conflict with naturalism and I would suggest that it lies at the root of various confusions in linguistic Platonism. In fact the analogy with physical objects is responsible for the defective type/token distinction presented by Platonists, specifically by forcing a singular denoting term reading of abstract objects analogous to that of physical objects.

If we persist in modelling the type/token distinction with this problematic definition of abstract objects as abstract physical objects, we will be stuck with an disjoint ontology and an epistemological conundrum as to how we can know the latter in the first place. *Ante rem* realism does not possess this particular drawback, among other things. For instance, if types are on the level of offices (in the sense discussed above) the analogy with physical objects is dropped, since these offices are not complete (do not have determinate truth values for all properties), do not have hidden natures and are certainly not sparse (offices can be created *ad infinitum* independently of entities discovered to fill those positions). For instance, for Millikan (2005) two semantic tokens are of the same type only if they are copied from the same pool of linguistic patterns or 'reproducing conventions' within a given community. Once the dogma of abstract objects is appropriately abandoned, the alternatives can be favourably illuminated. In the following subsections, I will discuss some of these alternatives and argue that at least one of them naturally dovetails with the existence of quasi-concrete structures as I have been describing them, specifically the last option.

10.4.1 Types as Archetypes

The following proposal, courtesy of Bromberger (1989), rejects the model of abstract physical objects discussed in the previous section but rescues a version of what is called the *Platonic Relationship Principle* in order to ground the type-token distinction in linguistics.

Bromberger valiantly attempts a reconciliation of three independently plausible views on the foundations of linguistics. (1) Linguistics is a theory of types, (2) information about these types are empirically grounded, i.e. in terms of physical tokens and (3) linguistic types are psychological in nature. The relationship between (1) and (2) directly concerns the type-token distinction while (3)'s connection to (1) is related, in my view, to wherein lies the intrinsic nature of our quasi-concrete objects. My focus will be on the relationship between (1) and (2) which Bromberger offers. I hope to show that it does not do justice to the type-token distinction we need.

On this view, (1) and (2) above are in principle grounded by the *Platonic Relationship Principle* or the principle that allows us to "impute properties to types after observing and judging some of their tokens" (Bromberger, 1989: 62). There are various elements of this proposal set within an interrogative framework and I will discuss the most important ones for the purposes of the current topic. Firstly, tokens on this view are modelled in terms of natural kinds or "quasi-natural kinds". To be a quasi-natural kind is to satisfy three conditions.

The first condition is a "modelling" condition, close in many ways to the picture I presented for grammars in the previous parts. To be a model here is to allow certain inferences from one random member of a class to properties of other members. It is a (dyadic) resemblance relation.

The notion of model on which I rely here is the notion of model appropriate for artefactual models. An artefactual model is an object so designed that, by finding the answers to some questions about it, a competent user can figure out the answer to questions about something else (Bromberger, 1989: 62).

The rhetoric then becomes familiar. A model of water can allow us to figure out facts about the number of hydrogen and oxygen atoms in a given water molecule. A map can tell us about the distance between two points in a real landscape from questions about the specific points on the map (and the scale of the map). In the parlance of previous chapters, by asking questions about the model, we can arrive at answers about the target system. He goes on to say "[n]o pair of objects stands (or fails to stand) in the model/modeled relation absolutely, but only relative to specific sets of questions, pairings of questions, and algorithms" (Bromberger, 1989: 63). The algorithm relates questions from the model to the modelled via a pairing relation such that not all answers to questions about the model are relevant to questions about the target (and it is closed under the "correctness" of answers). So, we can ask the same question and expect the same answer in terms of the boiling point of a given

sample of some chemical substance c and c in general (Bromberger has a convention of “Related Questions” to generalise this idea). The idea is that there are some questions aimed at the model which share an answer with the modelled such that there is no reason to distinguish the answer. There are of course questions which are not “projectible” in this way and pertain only to the model such as ‘What kind of paper is the map made of?’ or ‘In which test tube was the sample of chemical substance c kept?’ Projectible questions are thus the set of questions which receive the same answer for every member of the kind.

The next condition on quasi-natural kinds is the *explainable differences condition*. The condition states that differences among members of these kinds are to be systematic and law-governed.

So, for instance, samples of mercury differ in temperature, but the temperature of each at any given time is accounted for by laws and kinds of boundary conditions common to all samples of mercury at all times (Bromberger, 1989: 65).

This condition is couched in terms of the common presuppositions of the (nonempty) set of questions which receive slightly different answers for each member of the quasi-natural class. In other words, we need a way of tracking the differences of samples or tokens in a nomological and systematic way, what he calls “ w -projectible questions”. The cases in which these differences cannot be thusly tracked are then considered to be cases of presupposition failure. Asking for the boiling point of the word “mercury” results in a presupposition failure but samples of mercury all have w -projectible questions as to their boiling points on a specific occasion (but different answers to them).¹³

Lastly, we have the *individuation condition* which stands in contrast to the previous condition. The set of questions associated with this condition are those which have common presuppositions which are neither projectible nor w -projectible. For instance, one can ask of any word token when or where it was tokened or uttered (written etc.) but the different answers one would receive for different words would not be determined by law-like or systematic principles, they would be unrelated contingent matters.

With these three conditions in place, we can attempt to describe Bromberger’s type-token distinction. Tokens are members of a quasi-natural kind and types are archetypes of those kinds. Types are not the kinds themselves.

Instead he views the type as what he calls the archetype of the kind, defined as something that models all the tokens of a kind with respect to projectible questions but not something that admits of answers to individuating questions.

¹³Interestingly, Asher (2011) offers a very similar account in terms of presupposition failure of semantically anomalous sentences and category mistakes.

Thus for Bromberger the type is not the kind itself, but models all the tokens of the kind (Wetzel, 2014).

Thus, to be an archetype of a quasi-natural kind is to be a model (in the senses above) of only the projectible questions pertaining to that kind. Whatever it is that determines the “correct” answers to the projectible questions of the members of the kind also determines the archetype (the “determinables” of the questions as Bromberger puts it). But we have presupposition failure when either *w*-projectible or individuating questions are asked of archetypes such as “Where is the sentence ‘All men are mortal’?” or “Who tokened the sentence ‘All men are mortal’ on June 29th 2016?”

The *Platonic Relationship Principle* is justified in the case of quasi-natural kinds since answers to projectible questions are informed by the same physical empirical determinables (truth-makers?) directed at both the archetypes and the normal members of the quasi-natural kinds. Therefore, linguistics is indeed a theory of types (as Katz insisted), (1) and (2) above are perfectly compatible and the type-token distinction is vindicated.

Objections

This framework might seem compelling at first glance. However, upon further inspection it leads to untoward consequences and unclear resolutions.

Bromberger does do an admirable job of defending the intuition behind Platonic abstraction (or the Platonic relationship principle) without falling prey to the worries presented by Linsky and Zalta above of modelling abstract objects directly on the basis of ordinary objects. But there is an extent to which projectible questions do not fully describe the properties we often associate with types. Consider the first desideratum of a realist account of linguistic ontology advocated in sections 8.1 and 9.1 above, namely that we need to be able to account for the (potential) infinity of sentences of natural language. Questions of the cardinality of natural language and more importantly the answers to those questions will generally not be part of the projectible question set, since no token sentence or member of the quasi-natural kind of sentences will have a related answer. ‘How many centre embeddings are possible in this sentence type of English?’ is a very different question from ‘How many centre embeddings does this token sentence of English have?’. In addition, Wetzel (2014) claims that “generally there are no such properties had by all and only tokens of a type, at least in the case of words—not same phonological structure, nor same sense nor same spelling”. Some tokens just do not have Bromberger’s natural projectible properties. If the only route to the vindication of the Platonic relationship principle is *via* questions which have the same answers for both types and tokens, then formal properties or abstract rules of natural language types will not receive any characterisation. There seems to be an inference from tokens to types which is

missing (perhaps an induction step?). Another way to put the point is that natural languages have formal properties which will not receive any treatment if we employ the “bottom-up” Platonism of Bromberger.

In the opposite direction, we might have projectible questions which seem inappropriate to ask of types. Imagine a scenario in which all the members of a particular endangered species of wild animal, let’s call it a *jorra*, were in captivity. At some stage, in order to ensure the survival of the species, they are injected with a vaccine for a particular disease *d*. Let us further assume that their numbers are so few that it was possible to inject them all concurrently and in the same general location. In such a case, we would have a set of projectible questions such as “When were the *jorras* inoculated?” or “At what time was this *jorra* inoculated?” receiving the same answer for each individual animal in the species, let’s say “the 30th of June 2016 at 5pm”. In such cases, contingent facts about tokens seem to respect the *related questions convention* and thus pertain to the types or archetypes as well. But there is something amiss about this consequence for linguistic types. One could of course insist that these questions are not projectible or result in presupposition failure for archetypes but doing so would require a prior notion of archetype distinct from the mere interaction of tokens. Wetzel (2009) discusses a similar worry when she writes,

[T]here has to be much more to the story than “we know about a type on the basis of interaction with its tokens,” because there are uninstantiated types—some very long sentences, for example—about whose properties we can talk perfectly well (38).

The above concern is related to a larger one, namely that it is not easy to read off the ontological commitment of these questions and thus the ontological status of archetypes of quasi-natural kinds remains unclear. Bromberger’s account requires that the properties of types are exhausted by answers to the projectible questions of their tokens or quasi-natural kind. He states that “each archetype is characterized by a set of question-answer pairs, namely the set of projectible question-answer pairs of some quasi-natural kind, and is therefore exhaustively described” (1989: 69). But if the above arguments show anything, it is that there are properties of types which are not so characterised and thus require some additional characterisation, from the top-down if you will. It is unclear how to resolve the problems cited above without resorting to such a strategy and yet this tactic would surely be inimical to the proposal set out by Bromberger.¹⁴

¹⁴At this stage one might respond that (3) or psychological facts determine the answers to these projectible questions. Bromberger holds that certain classes of quasi-natural kinds form categories just in case they share common *w*-projectible and individuating questions but differ in their answers to projectible questions. Linguistic objects then share a category with psychological objects. But a Platonist could argue that linguistic types could similarly form a category with mathematical objects by the same reasoning.

Lastly, there is a problem with the definition of the modelling condition in terms of “correct” answers to questions. As we have seen, there are many different linguistic frameworks which generate the same sets of sentences (are weakly equivalent). Bromberger considers questions as to the S-structure (or tree form) of sentence types to be among the projectible questions of a given sentence type. Correctness seems to suggest that there is one unique structure which determines the answer to a question posed of its sentence type. In this thesis, I have urged against this picture for linguistics. This objection can be thought of as a variant of Quine’s problem of equivalent mental grammars with which we dealt in adopting the structuralist framework. We would not want to give that benefit up without a fight.

Correctness also militates against the possibility of tokens misrepresenting their types. It seems that there is no room for members of a quasi-natural kind failing to share projectible questions with the archetypes of that kind. Yet there are cases of tokens misrepresenting their types. Take the example of national borders (modified from Szabó (1999)). A given fence or wall could fail (due to contingent factors) to accurately represent the (abstract) border between two countries since certain questions as to its length could fail to project or have a different answer to the same questions posed of the abstract border. Similarly, individual uses of words could be out of sync with the types they are meant to exemplify (as in Szabo’s “inverted word argument”, see section 10.4.3 below).¹⁵ We would ideally like to be able to explain this phenomenon within our account of the relationship between types and tokens. Unfortunately, these and other issues make Bromberger’s approach unappealing.

10.4.2 Types as Wholes

So far we have seen a proposal in the spirit of a somewhat Platonic resolution to the type-token distinction (albeit bottom-up). In this section, I will review a radically different approach rooted firmly in physicalism. Kaplan’s project in ‘Words’ (1990) is to offer an alternative framework to the Platonic type-token distinction in terms of what he calls the *common currency* conception. Kaplan sets himself the target of word individuation but what applies to word types will presumably have significance for how we deal with other expression types such as sentences.

On the standard type-token distinction (exemplified by the quotations of Pierce and Quine at the beginning of this section), word types are physically instantiated by their tokens in the form of utterances and inscriptions of various sorts. Kaplan believes that this conception stems from the logician’s tool-kit and might be appropriate for algebra or a type-setter’s task but fails as an appropriate conception for a naturalistic account of words and language. He states that,

I propose a quite different model according to which utterances and inscrip-

¹⁵What of malapropisms? They seem to be tokens of word types which would receive different answers to projectible questions such as those pertaining to phonological structure.

tions are *stages* of words, which are the *continuants* made up of these interpersonal stages along with some more mysterious *intrapersonal* stages (Kaplan, 1990: 98).

One apparent immediate advantage of this new framework, over its Platonistic rival, is that it can account for the fact that tokens of types often do not share very many properties among themselves. Word tokens notoriously differ in their written and spoken forms. “Even the spoken tokens of a given word will exhibit a range of pronunciations due to accents” (Wetzel, 2009: 38). Writing conventions are similarly diffuse (although to a lesser extent). Surely we would want the two tokens “colour” and “color” to be related to the same type or the British and US pronunciations of the word *schedule*, ([ʃɛdju:l] versus [ˈskɛdʒul]). On Kaplan’s account the individuation conditions are such that these tokenings all refer to a single word in each case. This account also allows for the possibility that two distinct words can be in possession of the same meaning, spelling and pronunciation. I will return to that alleged phenomenon below. For now, we will survey some other features and advantages of adopting this particular metaphysical picture of words.

For one thing, the physicalism of this view seems to be more in tune with how words evolve in language than is Platonism. On the latter view, words have eternal, unchanging natures which are imperfectly reflected by their physical instantiations. Kaplan adds that the Forms of these words are represented by their spellings. I do not think that this link to orthography is either necessary or necessarily true. The conventional nature of orthography is generally appreciated and in certain cases, such as multiple transliterations of foreign alphabets, very apparent. In any event, on the stage-continuant conception, the history of a word, not its “phonographic” appearance, determines its identity. Hence the claim that distinct words can nevertheless have the same semantic, phonological and orthographic form since they are capable of having distinct histories.¹⁶

Another advantage of this approach is its focus on the human creator in the context of words. People after all are the creators of lexical items in various languages, through either naming practices or the like, and these items are then transmitted through interpersonal contact.¹⁷ The human element goes a step further in providing an analogy of word identity in terms of personal identity. People tend to exist over a period of time in which we undergo vast changes at both the cellular and psychological levels. Nevertheless, we remain the same person (aside from some metaphorical uses of the word) through that time period. Similarly, instead of accounting for word change in terms of a chain of replacements, “we can use the notion of a single entity

¹⁶This possibility relates to Kaplan’s other agenda of resolving certain puzzles in the philosophy of language, such as Kripke’s Paderewski case. I will not attempt to address this further agenda here.

¹⁷A quirky alternative might be to view human language users as transmitters or “vehicles” of words conceived of as genes in the rhetoric of the *Selfish Gene* (1976). Thus, natural languages would use us as tools for their survival over time.

undergoing *changes*" (Kaplan, 1990: 101). The story of word creation then proceeds in a similar fashion to Kripke's account of initial baptisms. There is an initial dubbing, then transmission in various media from person to person and linguistic community to linguistic community resulting in changes along the way. Words, like travelling salesmen, retain their identity throughout the vagaries and vicissitudes of their existence.

This is thus a part-whole conception of the metaphysics of words. Hawthorne and Lepore (2011) claim that it further entails a four-dimensionalist perspective. However, I tend to side with Bromberger (2011) in rejecting the need for such an ontology in Kaplan's case. One reason I have for siding with Bromberger is that Kaplan's view seems to require a distinctive temporal component. Unfortunately a discussion of the differences between 3D and 4D views or the A-series/B-series distinction in metaphysics are beyond the current scope. Either way, the view under discussion is physicalist in spirit. No part of the continuant and certainly none of its stages are to be found in an abstract realm of entities. Nor are Linsky and Zalta's worries of concern here, since types *qua* continuants are not modelled on the physical characteristics of their tokens or stages. The types are a conglomeration of the features of their tokens over time (or "archipelago" as Hawthorne and Lepore call it).

The last aspect of this view which I deem important for present purposes falls under the banner of the "mysterious" *intrapersonal* stages of words. A central question for Kaplan is "when does a word count as a repetition?" On the standard Platonic view, we could think of all tokens of a particular types as repetitions in a sense since they are all copies (albeit imperfect ones) of the type. Bromberger too can account for the repetition of a word if two or more members of a quasi-natural kind share both the answers to the projectible and individuating sets of questions. For Kaplan, matters are not that simple. For instance, how are different performances of a word connected to the same space-time worm given that its phonographic profile can be so distinct from other stages? Imagine a scenario in which you hear a speaker utter a particular word and then a moment later are expected to reproduce that utterance.

The identification of word uttered or inscribed with one heard or read is not a matter of resemblance between the two physical embodiments [...] Rather it is a matter of intrapersonal continuity, a matter of intention (1990: 102).

Kaplan insists that "[t]his notion of repetition is central to my conception" (1990: 103). Words are physical entities on his view, their transmission is a physical process and without a concept of repetitive communicative transmission, word survival would seem mysterious. Therefore there needs to be an account of how this "hand-me-down" process works. Kaplan's answer starts with what Hawthorne and Lepore call the "constitutive role of intention". In other words, what makes a word a repetition is that the speaker *intends* to create a repeat performance of that word. There is some

margin for error and dysfluency of imitation but the fact of the matter is settled by the speaker's intention.

Objections

There are a number of reasons that I think this picture won't do for the type-token distinction in linguistics. I will start with a few minor issues and then argue that they culminate into a significant problem for the view.

One initial problem is specifically related to the task I have set myself (and to which Kaplan perhaps is unconcerned). In order to interpret linguistic theory or its models at the appropriate level of types, we need a concept of an uninstantiated type. Linguists study sentences that have not been tokened as of yet, so these types have no initial baptism and thus no three-dimensional continuant in existence (they might have four-dimensional existence though). As Hawthorne and Lepore point out, this conception "assumes there are no unperformed words" (2011: 477). Their example involves derivational morphology and the possibility that some prefixed version of a word such as *unhappy* might never have been produced but it still a word nonetheless. Morphology allows for the introduction of many new words of a specific category with resources already available in the language.

This latter point brings us to Bromberger's central complaint with Kaplan's (and Hawthorne and Lepore's) framework, namely that the philosophical musings have transmogrified the real linguistic order of things. In a damning pronouncement, Bromberger puts the point in this way.

Their [Hawthorne and Lepore's] decisive objection is that Kaplan's "model" and principle of individuation require that each time we use a term, we use it with the intention of mimicking a specific previous use or memory icon of that term [...] and thus never access the generative powers of our morphology" (2011: 489)

He adds that if they are correct on this point (which he thinks they "indubitably" are) then there is nothing significant left of Kaplan's account, since it would only describe a small number of cases (if that). Bromberger's own objection is more in line with the structuralist project I have undertaken in the previous sections. He claims that Kaplan's view does not cover nearly enough ground to approach the ontology of words conceived of more functionally. A definitive aspect of word individuation concerns their roles in various hierarchical structures of syntactic theory, i.e. their roles as constituents in larger expressions. Without such an extension we cannot distinguish "John is eager to please" from "John is easy to please", in which John plays two distinct roles and similar syntactic data. Furthermore, I would add, that even simple adjectival phrases would not receive a treatment on Kaplan's isolationist approach. There is a difference between a red apple, a red pen and someone who has red hair.

Bromberger's further objection is that not only is Kaplan's account impoverished as it currently stands but it also "distorts and diminishes" lexical mastery of a language so much so that it is rendered incompatible with any legitimate account of that mastery.¹⁸

The above considerations point to the central failure of this view, it is unclear how to generalise it to other expression types at all. Are sentences greater wholes of which words are parts? If so, what of the combinatorics of these expressions, à la Bromberger? How does repetition and intention play a role in the metaphysics of sentences?

The target of Kaplan's project is so *sui generis* that it at best cannot be used as a launching-pad for the investigation of other quasi-concrete objects or the structures in which the units of our languages are set and at worst fails to provide an account of the full nature of words themselves.

10.4.3 Types and the Representation Relation

At this point, I would like to draw attention to an underlying motivation behind this detour into the ontology of words and the type-token distinction. The ante rem realism advocated in previous sections was developed with the rejection of Platonism (and nominalism) in mind as the best means of capturing the requisite realist view. This view involved the modification of pure (ante rem) structuralism to include the existence of quasi-natural objects or structures so infused according to Parsons (1990).

To review, quasi-concrete objects, such as words, are those objects which have more than just relational properties capable of characterisation in purely structural terms (such as sets of axioms).¹⁹ They also possess intrinsic properties or what Parsons calls either an instantiation relation or a representation relation in the concrete. In this section, I want to make the claim that it is not an arbitrary choice as to which relation constitutes these objects for the particular ante rem proposal under discussion. The reason for this claim is essentially rooted in an old debate about the existence of universals staged between Plato and Aristotle.²⁰ More precisely, if we conceive of quasi-concrete objects in terms of an instantiation relation we smuggle a version of Platonism back into the framework since instantiation presupposes a Platonic conception of universals. Needless to say, this would be a most unfortunate

¹⁸On a related note, Bromberger criticises both Kaplan and Hawthorne and Lepore for offering too parochial of an account of word change. He argues that the conception of word change "is at best a shortcut, and a way of being noncommittal about empirical details" (Bromberger, 2011: 497). Language change involves factors ranging from the sociological and psychological to the political and fashionable and tends not to be focused on the changes involved in individual words *in vacuo*.

¹⁹I also suggested that Devitt's RESPECT might be a means of tracking our choice of relation for linguistic objects.

²⁰I thank Zoltan Szabó for drawing my attention to this general concern and specific point.

result tantamount to inconsistency. Hence the need to banish this dogma. It is therefore necessary to find another route to identifying the nature of linguistic objects. I will propose a thesis to the effect that the relationship between types and tokens necessary for the definition of quasi-concrete objects is to be defined in accordance with a proposal made in Szabó (1999) in terms of the representational nature of the aforementioned relationship.

Szabó pitches his framework in direct contrast to frameworks which endorse “the instantiation view” or “[a] type *T* is instantiated by its tokens, and it is in virtue of this that empirical information about a token of *T* can play a role in justifying our knowledge about *T*” (1999: 147). If this view sounds familiar it is because it closely resembles Bromberger’s *Platonic Relationship Principle*. Szabó’s characterisation goes further to insinuate that most views which incorporate types as special kinds of sets or patterns with their tokens exemplifying a certain kind of (projectable) similarity are implicitly committed to the instantiation view.²¹ Kaplan’s account above is a notable exception.

We have already seen some problems with the traditional type-token distinction. Generally, types are identified by features of the tokens, either the way they sound or are spelt (what Kaplan calls “phonographic” features). However, these features are generally unreliable sources for getting at the nature of types. Szabó follows Kaplan, Wetzel and others in pointing out that “recognitional criteria” (based on the “physical appearance” of tokens) are often (although not always) unhelpful in the pursuit of type identification.

Categorizing tokens like this would make types linguistically widely heterogeneous, in a way that would imperil the reliability of inductive inferences from tokens to types. Even if phonological and orthographic criteria are acceptable for some purposes, they are unacceptable for explaining our linguistic knowledge of types (Szabó, 1999: 148).

The point is that the forms or physical appearance of words and sentences and other tokens are part of a motley assortment of tools we use for the identification of types (which could include Kaplan’s intentions, semantic information, context and even linguistic theory). Proponents of the instantiation view tend to neglect this fact. I think that despite his disavowal of the view, Kaplan too is culpable of throwing the baby out with the bathwater in this way. It is important to note that Szabó claim is not that the instantiation view is untenable, quite the contrary, he believes it in principle to be “simple and plausible” but he offers an alternative which aims to better capture the nature of the relationship between types and tokens. It is on to this proposal or “representational view” that we now move.

²¹This is the main reason I did not consider Hawthorne and Lepore’s *abstracta-articulations* model in any detail. The assumption being that general objections to the instantiation view will militate equally against their *representation* of it.

A good starting point for the appreciation of the difference between the representational view and its Platonic rival, is an insight which dates back to Aristotle. ‘Types are nothing more than abstract particulars’. The central idea of this claim is that types are incapable of instantiation or they cannot have instances. In Ancient Greek syntax, attaching the definite article to any noun (or infinitive, if the neuter article is used)²² results in an abstract version or concept pertaining to the referent of that word. For instance, from “good things” (one word in Greek) we derive “the good” (ἀγαθοί → τὸ ἀγαθόν), from the infinitive “to do wrong” to “the wrong” or simply “injustice” (ἀδικεῖν → τὸ ἀδικεῖν). In English, singular terms often fulfil this role. “We talk about the first line of Gray’s *Elegy*, the last words of Goethe, of the fourth letter of the Hebrew alphabet” (Szabó, 1999: 152). In these cases, words, infinitives and singular terms are being *used* to fulfil a certain role, namely a representational one.

The sort of reasoning employed in the previous paragraph indicates a functional approach to the role of tokens and their relationship to types. The function of a token is to represent a type and it is common ground that they play this role in the language. In this way, tokens are representations which stand “proxy” for what they represent. There is an element of arbitrariness to this relation. Unlike the idea that tokens are instances of an overarching type, which suggests a more intrinsic picture of the role of tokens, on this view tokens are related to the class of objects which have Grice’s non-natural meaning. In this sense, Bromberger’s initial “modelling condition” wasn’t that far off in that it emphasised the multiple realisability of tokens *qua* models. He did not, however, push the artefactual nature of tokens far enough.

The story goes on to specify exactly what kind of representation relation tokens bear to their types, namely an indirect one.²³ By way of example, Szabó writes that “[t]he English word-type ‘horse’ represents horses and so do all its tokens. But unlike the word-type, those tokens represent only *indirectly*: they represent the word-type ‘horse’, which in turn represents horses” (1999: 150). Thus, this view encompasses two stages from tokens to the referents of the types, we learn about the nature of types not only by looking at the features of the tokens but also by knowing their roles as representations. To use a variant of Szabó’s example, if we want to know about a particular species of animal, we use not only reference to that animal but also reference to representations of it (as found in books or pictures). Tokens are then in this sense very much like models in scientific discourse. They act as intermediaries between the types and the kinds of things that types represent. The important difference between this view and the instantiation account is that instances require additional ontological connections with their types that representations do not. This is why a sign in British Sign Language, a spoken speech act in Cockney English, Morse code and a written word can all represent the same word type without

²²German too shares this feature as far as I am aware.

²³Here Szabó’s account is analogous to my view of semantic modelling in Part II.6.

issue. To find the set of shared features through which they instantiate that type is a more daunting task than to understand their roles as representations.²⁴

This view allows for a lot of flexibility which in turn can explain more uses of tokens. The problem, Szabó acknowledges, is that many of the reasons for and against this proposal can seem to lead to an irresolvable conclusion (as with many of the original debates concerning universals). What is a natural explanation for a representationalist might be unnatural for an instantiationist. In light of this, Szabó has a specific argument based on inverted spectrum arguments, tailored to bring out the advantages of the former position. I will summarise the idea below.

Imagine a scenario in which a second language English learner inverts the mapping between numbers and numerals such that certain teens (i.e. from 13 to 19) are systematically confused for their counterparts in the terms of multiples of 10 (i.e. from 30 to 90). Now imagine further that this speaker performs the inverse mapping in speech. So whereas she might write '14' to indicate '40', she would not pronounce the mistake. When asked about her age (which we can assume lies within either range), the mistake is not detectable and the utterance seems to be true. But there is a problem here. By some Gettier-like luck, the speaker happens to make what seem to be true numeral utterances when prompted in speech, but in writing the errors are exposed. The problem is that the speaker/learner makes a mistake about the representational value of certain tokens (numerals). If tokens are merely instances, it would be harder to explain the distinctive error in such cases. All the speaker would be doing erroneously, would be failing to categorise the token of say '40' appropriately in the same way that believing '40' to be a verb or a mass noun would be a failure of categorisation. The two-part representational analysis explains why the speaker lacks the knowledge of the referent whereas the instantiation view does not. "The inverted word argument shows that knowledge of reference is mediated by knowledge of tokening" (Szabó, 1999: 155).

The point of the above argument is that in order to know the reference of a term, one needs to know the types represented by the tokens. And as we have seen (also with the national border case), tokens can fail to represent their types and we can make errors, an account of which the instantiation view seems inadequate to provide. In its two-part representational picture, Szabó's view allows for a more nuanced account of the error in the inverted word scenario in terms of tokens failing to represent the appropriate types and this failure impeding knowledge of reference.

There is a profound worry within this proposal. The worry is that despite all the best efforts to distinguish the representational view from the instantiation view, when confronted with ontological questions, the former allows for a version of Platonism to resurface like a dormant virus. Furthermore, the ontological package with which

²⁴Presumably there are two notions of representation at work here. Tokens represent types in a different manner to which types represent their referents. For instance, tokens might "stand proxy" for types but it doesn't seem to be the case that types are proxies for their referents. The type *horse* does not stand proxy for horses.

this view comes is supposed to be purely physicalist one. Types or those things which are represented by tokens are not supposed to be ontologically occult, they are not Platonic forms. Certainly, there is nothing inconsistent about supposing that the *representata* are physical in some sense but exactly what this sense is is not an easy matter to resolve. This is where Szabó's account becomes (self) admittedly "tentative and speculative". I think that it is at this point that the representational view can be buttressed by the ante rem realist framework presented for linguistics in this thesis. I will offer some details in that vein, among other things, in the following subsection. But first I will describe some advantages of incorporating the representational type-token distinction into the larger ante rem realist picture.

Advantages of this view for Ante Rem Realism (and *vice versa*)

Besides the fact that this view rescues ante rem realism from falling back into Platonism, its adoption has a number of other benefits. I will mention the most important one below before considering how ante rem realism might assist the representational view in a similar escape from a Platonic ontology.

The original objection which Parsons posed to pure structuralism was that there are certain kinds of objects which cannot be described from a purely external relational perspective. The reason that these objects are recalcitrant to such characterisation is that they possess an "intrinsic" representation or instantiation relation in the concrete, hence the name "quasi-concrete". Thus, the claim is that *there is something about* tokens such as inscriptions on a page etc. which they possess, rooted in the concrete, in addition to their relational properties. We have already seen the perils of opting for the Platonic interpretation of quasi-concrete objects in terms of instantiation in section 10.4.1 and above. The option I proposed, with Szabó (1999), was that this "something" concrete about quasi-concrete objects is their representational capacity. Tokens of linguistic types, such as words or sentences, possess an ability to represent types of all kinds outside of their structural properties.

In a sense, this capacity is "intrinsic" as Parsons puts it, in that it is part of the role of a linguistic token that it represents a type in addition to featuring in larger structural configurations. However, I would be cautious to construe this intrinsic nature in essentialist terms. Tokens have a non-natural ability to represent types, which is why tokens of varying phonographic profiles can represent a single type and tokens of a similar or the same profile can represent distinct types. There is a certain arbitrariness to this relation. Tokens of all sorts can perform (and fail to perform) in their representational tasks but they are generally defined by that capacity nonetheless. Following Szabó's picture, word tokens and sentence tokens etc. represent types and indirectly referents. Departing from this picture (but hopefully still within its spirit), the types do not refer to or represent objects in the world directly (or abstract objects) but represent nodes in larger linguistic structures in either the places-as-objects or the places-as-offices perspective (see section 10.2 above).

Ante rem realism was developed as an ontological alternative to Platonism. It is at this point that the view might be of use to the representational view of tokens and their types. Szabó has a distinct account of what types are not, i.e. Platonic forms. His view of what they are, ontologically speaking, is more vague. He holds that types are created, not discovered, and thus that they have starting points in time. The problem with this claim is that it requires a certain creative power to be attributed to representation, since it presupposes that there must have been a first token of any given type. Without ever-existing Platonic types of words, this is hard to maintain. In the following I will argue that ante rem realism can buttress the representational account such that it neither needs Platonic forms nor the explicit moves made by Szabó in order to avoid them.

In terms of ontology, Szabó makes a few ingenious observations but also concedes too much in others. In terms of the former observations, he proffers a more nuanced picture of representation such that it can occur before and independently of the existence of a *representatum*. On the intuitive copy-model of representation, word-types must exist before they can be represented by tokens, as an original must exist prior to its copies. This latter model, however, cannot account for some cases of representation such as the one below.

Most of the work of an architect consists in producing representations –floor plans, drawings, models, detailed descriptions – of buildings that do not yet exist. Once we abandon the copy-model of representation, there is no difficulty here. The representation view can coherently maintain that the first tokens of a new word-type are much like the drawings of the architect: they represent something that does not yet exist (Szabó, 1999: 162).

I think that there is something to this idea but it is too quick in its current form. An adherent of the copy-model could object that the architect does not produce representations of non-existent buildings but rather uses existing representations to assemble novel structures. In this way, the creativity of representation resembles the creativity of language-use in terms of compositionality advocated in section 8.1 and 9.1 earlier in this chapter. At the sentential level, Szabó's idea is sound. There are types of sentences determined by the rules of language, much like there are types of buildings determined by the laws of design and physics, that have yet to be tokened. The architect's blueprints represent something more abstract than a particular building, they represent a building structure-type. The point here is that attempting to describe representations *sans* the *representatum* is problematic if construed in a piecemeal manner but with relation to systems of structures it might be more appealing. Nor do the structures themselves need to exist *in toto* for the process to work. As the rules of the language develop and evolve, so do the systems of rules which represent the yet-to-be represented structures of that language.

The copy-model also incorporates an independence claim or the claim that the thing represented must exist independently of its representations. The necessity of

this condition on representation is also challenged by Szabó by means of the national borders case. There are various ways of representing a national border (fences, walls, xenophobic attitudes etc). If these ways of representing are eliminated, so too are the borders, Szabó argues. Thus, the border does not exist independently of the objects which represent it. Indeed, a national border seems like a good candidate for a quasi-concrete object. But I wonder if this reasoning applies to objects like the equator. There are ways in which we can represent the equator, imaginary lines across the circumference of the planet (or real lines drawn on representations of the planet), perhaps a long physical tube across the surface of the earth etc. However, if these representations are destroyed the equator remains. One could think of a national border analogously.²⁵ Of course, destroy the planet and with it go the equator and all national borders but that is beside the point for now. The problem with the example is that one could ask how the representations of the borders are set up in the first place if there were no independently existing templates from which to work. For this reason, I think that these examples fall into the camp of bad prospects for a “quick resolution” which Szabó dismisses earlier in his paper. Worse still, they force him to concede that certain linguistic types “cannot exist untokened” (Szabó, 1999: 162). Furthermore, this possibility then produces a distinction in kind between the complex expression types (e.g. clauses and sentences) and the simple expression types (e.g. morphemes and word roots) of which these types are comprised. He offers a Kaplanian suggestion for how the histories of word-types might have transpired to lend credence to this new distinction.²⁶

I argue that on a more structuralist account, such as ante rem realism, the concessions made above, that tokens can represent in the absence of types and that words and sentences are representationally distinct, are not strictly necessary. In terms of the latter, following Kaplan on word histories and the like might commit us the same error which Bromberger (2011) points out, namely that we are starting from a narrow conception of word-change over time as opposed to a more empirically sound conception of language change over time. To recap, Bromberger urged two separate

²⁵There are differences between the two objects. For instance, the equator is a mathematical feature of an object shaped in a certain way. One could argue that its existence is necessary or at least necessarily dependent on the existence of that entire object. National borders, on the other hand, are political objects defined by treaties and historical contingencies. Nevertheless, once they are so conceived they can take on mathematical properties.

²⁶Hawthorne and Lepore’s objection to Kaplan also militates against this suggestion in terms of the generative capacity of derivational morphology. To take this idea a bit further, the line between word and sentence seems to be quite clear in *isolating* languages such as English or Afrikaans but this is not the case for certain members of *agglutinative* and *polysynthetic* language families, where what we would render with an entire sentence can be expressed with a single unit (with affixes appropriately appended). For example, in some of the Yupik languages of parts of Alaska and Siberia larger expression types typically involve a one word root upon which various suffixes are added in order to create sentences or sentence-like structures. Consider an example from Steven Jacobson (1984) of the Central Yupik word *angyaliurvigpaliciquq*, which can be translated as “he will build a big place for working on boats.”

points. The first was that a proper account of word-type needs to account for the structural elements of words, i.e. “that words function as constituents of phrases and sentences”. He took this to be the essence of words, that they function in larger structures “whatever their intrinsic perceptual and referential features”. The point might be overstated and I think that a proper account of words should consider both representational and structural components of words in terms of their quasi-concrete natures. The second point was that the focus on initial dubbings and individual word change is a red herring. In other words, Kaplan’s attempt to anthropomorphise words is a non-starter. In fact, the sociological perspective is a better place to start than the individualistic one.

Normally, single words do not change in isolation, but whole families of words that share features change together as certain shared constituent features get replaced in shared phonological environments (Bromberger, 2011: 497).

He goes on to cite examples from the Great Vowel Shift in English to the Valley Girls Rise in North American dialects. I think that this line of objection might be damning for Kaplan’s (and Hawthorne and Lepore’s) view but not so for Szabó’s. The representational view buttressed with an *ante rem* realist account of quasi-concrete objects offers a route to accommodating both of Bromberger’s worries. Firstly, structural elements of words are readily accounted for in this framework since it is at base a structuralist view of linguistic objects. Secondly, any changes are naturally pitched at the level of structures and substructures through which natural languages are characterised. *Ante rem* realism is thus in tune with linguistic literature on language change which is often described as law-like (as Bromberger asserts) and integrated, in that changes have structural effects and do not generally occur in isolation with individual histories. It is therefore not guilty of the linguistic *façons de parler* Bromberger so strongly opposes for the ontology of words. Nor do we have to drive a wedge between our ontological treatment of words-types and more complex expression types such as sentences since both words and sentences are quasi-concrete construed in terms of the representational view of types and tokens.

10.5 Subconclusion

Since the late 1950’s linguists have discussed linguistic structures, their implementation in grammar formalisms and their interrelations. Very little has been said specifically about what these structures are and how they relate to other non-linguistic structures. I have attempted to give the beginnings of an account here. Much work still needs to be done. Nevertheless, *ante rem* realism provides not only an account of the foundations of linguistics and its subject matter but also aims to demystify the concept of structure used throughout the discipline as an abstract pattern produced by competence but distinct from it in ontology.

The question remains, what kind of science is linguistics? Is it a formal science in terms of mathematics or an empirical science like psychology. On the view I have been pushing, the answer is that it is a little bit of both. One could either take it to be an empirical science with formal aspects or a formal science with empirical aspects (depending on your funding grant), it lies in the same disciplinary lacuna that most applied sciences do (see Part II of this thesis).

In this chapter, I have attempted to provide a more sound realist footing for the foundations of linguistics than is provided by the traditional Platonism of Katz and Postal, in the hopes of offering ontological support to the structural realist suggestion of the previous part. I argued for three conditions or desiderata on any realist account of linguistic ontology in light of critiques found in the Platonist and nominalist literature, namely creativity and infinity, the correct level of scientific abstraction and RESPECT between the distinct structures of the mind and linguistic world respectively. I then showed that Platonism cannot meet these conditions.

Lastly, I drew from the philosophy of mathematics to provide a novel account of the nature of the linguistic enterprise and the natural languages it studies, in terms of an *ante rem* or non-eliminative structuralism with the inclusion of quasi-concrete structures, which I called *ante rem* realism. Within the course of developing this view, I took a necessary detour into the ontology of abstract objects with specific emphasis on the class of such objects to which I claimed linguistic structures belong, namely quasi-concrete objects. I defined them in terms of a representation as opposed to an instantiation relation in the concrete so as to avoid a resurgence of Platonism. This account aimed to meet all of the desiderata of a realist linguistic account in a way more amenable to naturalism.

Conclusion

One of the aims of this thesis has been to unite theoretical linguistics, the philosophy of science (particularly modelling) and the ontology of language. Each part of the research presented here targeted these goals separately with the unified aim of bringing greater clarity to the foundations of linguistics from a philosophical perspective.

In the first part of this thesis, I argued for a grade of mathematical involvement for linguistic grammars which was neutral on issues of ontology and the exact mechanisms underlying human language cognition. I described and condemned the mathematisation of natural language which draws untoward connections with arithmetic and in turn linguistics and the formal sciences. I went on to argue that further commitments as to the ontological or representational nature of linguistic grammars was deeply problematic (although not altogether hopeless). This reframing and reconceptualising of the traditional debate on the foundations of linguistics told a story of continuity between Conceptualism and Platonism in terms of the interpretation of the grammars and mathematisation of the object of inquiry. Nevertheless, at the level of each grade different positions on methodology were teased apart from views on ontology with the result that Platonism was not entailed by the third grade of involvement and Conceptualism could be consistently recaptured by the first. I consider these to be significant results in a debate which has not given much way in 50 years.

The task of the Part II was related to the philosophical agenda set forth in the first grade of involvement. Where the first grade and first part provided philosophical arguments for the rejection of various ontological constraints on grammars and also the view that grammars are scientific theories, the second part offered a descriptive basis in the actual practice of syntax and semantics with the aim of bolstering the position that grammars are in fact more akin to scientific models. This latter goal was achieved through two distinct strategies. In the first chapters, I surveyed the development of generative grammar with special attention paid to the modelling strategies employed in its service. I argued that two idealisation techniques were of particular significance to the linguistic project, minimal generation and isolation idealisation, both indicative of the minimalist idealisation of Weisberg (2007b, 2013). I argued that in linguistics, Cartwright's *simulacrum* account of the mendacity of

fundamental laws was particularly apt (perhaps more so than in physics). I then went on to recast the generative tradition purely in terms of modelling. This construal (following Blutner's (2011) similar construal on the theoretical side) allowed for a previously unconsidered perspective on the contribution of generative linguistics in a time when most of its core theoretical tenets are being challenged. This view also allowed for a novel connection to be made between generative grammar, optimality theory, the parallel architecture and dynamic syntax. The corollary of this connection was that the alternative models of model-theoretic syntax could be distinguished despite the often confusing discussions of weak generative capacity and notational variance.

Part II also saw a development of modelling techniques in formal semantics. In this chapter my aim was not only to survey the landscape with a novel interpretation in mind but also to offer an appropriate analogy between formal semantics and applied mathematics. In so doing, I provided a lens through which to appreciate the explanations present in the field and the syntax-semantics interface in a genuinely new light, as embedded indirect representation. Lastly, I briefly attempted to steer the discussion of modelling in linguistics away from the Scylla of instrumentalism and the Charybdis of anti-realism by suggesting a structural realist alternative in which the models of syntax and semantics provide genuine insights into real structures of natural language *via* indirect representation of those structures (in a way germane to the discussion in Part I of the first grade of involvement).

Part III investigated the ontology of natural language. It started by taking Realism seriously as an alternative to the Conceptualism (and Platonism) of Part I. Another way to think of this Part is that it aimed to support the realism of Structural Realism by offering an account of what linguistic structures might in fact be. It challenged and eventually rejected Platonism as a viable Realist theory of natural language and proffered a unique account in terms of what I called *ante rem realism*. This account drew from the mathematical structuralism of Shapiro and Resnik while modifying and amending their view to the specific exigencies of the linguistic enterprise, i.e. its empirical and *a posteriori* nature. One way in which I did this was through adapting Devitt's (2006) RESPECT relation connecting competence to the output of the grammar. I took a stance on the ontology of words and the type-token distinction in metaphysics and sided with a representational account proposed by Szabó (1999) while introducing some specific additions in light of the general framework presented in this part of the thesis. At the end of the day, I hope to stand this view up next to its various competitors in Conceptualism, Platonism, Pluralism and Inferentialism and argue for its advantages over these other frameworks for the foundations and ontology of linguistics.

Despite the generally ambitious task set for each part of this thesis, it holds many suggestions for further research and inquiry. For instance, the structural realist interpretation of the science of linguistics suggested in Part II is in need of development the likes of which were not possible in the present work (due in part to

focus and in part to space). Another interesting point of possible later investigation is the further connection between idealisations and abstractions in linguistics compared with the other social sciences such as economics, sociology and psychology. This strikes me as a particularly fruitful endeavour.

In the present work, I have attempted to provide a unified theory of natural language and the science which sets itself the goal of taming it. I have identified systematic errors such as mathematisation and offered novel construals such as the embedded indirect modelling of formal semantics. I have divorced methodology from ontology in Part I and set up a cutting edge framework for bringing them back together in Parts II and III. My hope is that this work can serve to cultivate the interest of philosophers of all persuasions into the appreciation of the genuinely rich theoretical terrain that is linguistics, its mathematics, models and structures.

References

- Asher, N. 2011. *Lexical Semantics in Context: A Web of Words*. Cambridge University Press.
- Awodey, S. 1996: 'Structure in Mathematics and Logic. A Categorical Perspective'. *Philosophia Mathematica* (3) 4. 209-237.
- Azzouni, J. 2004a. *Deflating Existential Consequence: A Case for Nominalism*. New York: Oxford University Press.
- Azzouni, J. 2004b. 'The Derivation-Indicator View of Mathematical Practice'. *Philosophia Mathematica* (III) Vol. 12, pp. 81-105.
- Azzouni, J. 2005. 'How to Nominalize Formalism'. *Philosophia Mathematica* (III) Vol. 13, pp. 135-159.
- Baggio, G., van Lambalgen, M., and Hagoort, P. 2012. 'Language, Linguistics and Cognition'. In Kempson, R., Fernando, T., and Asher, N. (eds). *Philosophy of Linguistics*. Elsevier B.V. North Holland.
- Bar-Hillel, Y. 1953. 'On Recursive Definitions in Empirical Science'. *11th Natural Congress of Philosophy* 5: 160-165.
- Behme, C. 2015. 'Is the ontology of biolinguistics coherent?'. *Language Sciences* 47, 32-42.
- Benacerraf, P. 1965. 'What Numbers Could Not Be'. *Philosophical Review*, 74. 47-73.
- Benacerraf, P. 1973. 'Mathematical Truth'. *The Journal of Philosophy*, Vol. 70, No. 19. pp 661-679.
- Biberauer, T., and Zeijlstra, H. 2012. 'Negative Concord in Afrikaans: filling a typological gap'. *Journal of Semantics*, 29: 345-371.
- Binder, J., Frost, J., Hammeke, T. Cox, R., Rao, S., and Preito, T. 1997. 'Human Brain Language Areas Identified by Functional Magnetic Resonance Imaging'. *The Journal of Neuroscience*, 17(1):pp 353-362.
- Bloomfield, L. 1936: 'Language or Ideas?' *Language*, vol. 12, pp 89-95.
- Blutner, R. 2000. 'Some aspects of optimality in natural language interpretation'. *Journal of Semantics*, 17: 189-216.
- Blutner, R. 2011. 'Taking a Broader View. Abstraction and Idealization'. *Theoretical Linguistics*. Vol. 37.1-2.
- Boeckx, C. 2015. 'Beyond Humboldt's problem: reflections on biolinguistics and its relation to generative grammar'. *Language Sciences*, Vol. 50, pp 127-132.
- Boolos, G. 2000. 'Must we believe in Set Theory'. In Sher, G., and Tieszen, R. (eds.) *Between Logic and Intuition: Essays in Honor of Charles Parsons*. Cambridge University Press.
- Bos, J. 1995. 'Predicate Logic Unplugged'. In *Proceedings of the 10th Amsterdam Colloquium*, Amsterdam: ILLC, pp. 133-142.

- Brandom, R. 2007. 'Inferentialism and Some of Its Challenges'. *Philosophy and Phenomenological Research*, Vol. 74, No. 3, pp 651–676.
- Bromberger, S. 1989. 'Types and Tokens in Linguistics'. In *Reflections on Chomsky* George, A. (ed.), pp 58–90. Basil Blackwell.
- Bromberger, S. 2011. 'What are Words?'. *The Journal of Philosophy*, 108(9): 486–503.
- Bueno, O., French, S., and Ladyman, J. 2002. 'On Representing the Relationship between the Mathematical and the Empirical'. *Philosophy of Science* 69: 497–518.
- Bueno, O., and Colyvan, M. 2011. 'An Inferential Conception of the Application of Mathematics'. *Nous*, 45:2. 345–374.
- Burten-Roberts, N., and Poole, G. 2006. "'Virtual conceptual necessity', feature-dissociation and the Saussurian legacy in generative grammar'. *Journal of Linguistics*, 42, pp 575–628.
- Bunt, H. 2007. 'The Semantics of Semantic Annotation'. In *Proceedings of the 21st Pacific Asia Conference on Language, Information and Computation (PACLIC21)*, pages 13–28.
- Bunt, H., and Muskens, R. 1999. 'Computational Semantics'. In Bunt, H., and Muskens, R. (eds) *Computing Meaning, vol.1*, Dordrecht: Kluwer, pp. 1–32.
- Büchi, J. 1960. 'Weak second-order arithmetic and finite automata'. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 6:66–92.
- Cann, R. 2001. 'Case: a dynamic approach'. Unpublished manuscript.
- Cann, R., Kempson, R., and Marten, L. 2005. *The Dynamics of Language: an introduction*. Elsevier Academic Publishers.
- Cann, R., Kempson, R., and Wedgwood, D. 2012. 'Representationalism and Linguistic Knowledge'. In Kempson, R., Fernando, T., and Asher, N. (eds). *Philosophy of Linguistics*. Elsevier B.V. North Holland
- Carnie, A. 2013. *Syntax: A Generative Approach. Third Addition*. Wiley-Blackwell Publishers.
- Carr, P. 1990. *Linguistic Realities*. Cambridge University Press.
- Cartwright, N. 1983. *How the Laws of Physics Lie*. Clarendon Press.
- Cartwright, N. Shomar, T and Suárez. M. 1995. 'The Tool-box of Science', in *Herfel* 1995, 137–150.
- Chomsky, N. 1956b. 'Three models for the description of language'. *IRE Transactions on Information Theory* IT-2:113–123.
- Chomsky, N. 1957. *Syntactic Structures*. Mouton Press: the Hague.
- Chomsky, N. 1959. 'On certain formal properties of grammars. *Information and Control* 2(2):137–167.
- Chomsky, N. 1965. *Aspects of a Theory of Syntax*. MIT Press, Cambridge, Massachusetts.
- Chomsky, N. 1966. *Cartesian linguistics: a chapter in the history of rationalist thought*. New York: Harper & Row. (2009) Third Edition.
- Chomsky, N. 1972. *Studies on Semantics in Generative Grammar*. The Hague: Mouton.
- Chomsky, N. 1980a. *Rules and Representation*. New York: Columbia University Press.
- Chomsky, N. 1981. *Lectures on Government and Binding*. The Hague: De Gruyter Mouton.
- Chomsky, N. 1982. 'A Note on the Creative Aspect of Language Use'. *The Philosophical Review* XCI, No.3. pp 423–434.
- Chomsky, N. 1986a. *Knowledge of Language: Its Nature, Origin, and Use*. New York Praeger.

- Chomsky, N. 1991a. 'Linguistics and cognitive science: Problems and mysteries'. In: A. Kasher (ed.), *The Chomskyan turn*, 26–53. Oxford: Blackwell.
- Chomsky, N. 1995. *The Minimalist Program*. Cambridge, Mass.: MIT Press. Reprinted 20th Anniversary Addition, 2015.
- Chomsky, N. (2000a). 'Minimalist Inquiries'. In *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*, (eds.) Martin, R., Michaels, D., and Uriagereka, J. 89–155. Cambridge, Mass: MIT Press.
- Chomsky, N. 2000b. *The Architecture of Language*. New Delhi: Oxford University Press.
- Chomsky, N., Hauser, M., and Fitch, W. 2002. 'The Faculty of Language: What is it, who has it, and how did it evolve?' *Science*, Vol. 298. pp 1569–1579.
- Chomsky, N. 2004a. 'Beyond explanatory adequacy'. In Belletti, A (ed.), *Structures and beyond. The cartography of syntactic structures Volume 3*, pp 104–131. Oxford University Press.
- Chomsky, N. 2005. 'On Phases'. MIT MS.
- Chomsky N. 2006. *Language and Mind*. Cambridge, UK: Cambridge University Press.
- Chomsky, N. 2008b. 'On Phases'. In Freidin, R., Peregrín Otero, C., and Zubizarreta, M. *Foundational Issues in Linguistic Theory. Essays in Honor of Jean-Roger Vergnaud*, 133–166. Cambridge, MA: MIT Press.
- Chomsky, N. 2010. 'Some simple evo-devo theses: how true might they be for language?' In R.K. Larson, H. Yamakido, and V. Deprez (eds.), *Evolution of Human Language: Biolinguistic Perspectives*. Cambridge University Press, pp 45–62.
- Chomsky, N. 2012. *The Science of Language: Interviews with James McGilvray*. Cambridge University Press, Cambridge.
- Clark, H. 1996. *Using Language*. Cambridge University Press.
- Collins, J. 2007. 'Review of Devitt 2006a'. *Mind*, 116, 416–23.
- Collins, J. 2008a. 'Knowledge of Language Redux.' *Croatian Journal of Philosophy*, 8, 3–43.
- Collins, J. 2008b. 'On the Linguistic Conception of Linguistics.' Manuscript.
- Collins, J. 2008c. 'A Note on Conventions and Unvoiced Syntax'. *Croatian Journal of Philosophy* 8, 23.
- Colyvan, M. 2001. 'The Miracle of Applied Mathematics'. *Synthese* 127: 265–277
- Cowart, W. 1989b. 'Notes on the Biology of Syntactic Processing'. *Journal of Psycholinguistic Research* 18(1): pp 89–103.
- Culbertson, J, and Adger, D. 2014. 'Language learners privilege structured meaning over surface frequency'. In *Proceedings of the National Academy of Sciences of the United States of America - PNAS*. 111, 16, p. 5842–5847.
- Culbertson, J., and Newport, E. 2015. 'Harmonic biases in child learners: In support of language universals'. In *Cognition*, 139, pp. 71–82.
- Davidson, D., 1965. 'Theories of Meaning and Learnable Languages.' Reprinted in *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press, 2001, 3 – 16.
- Davidson, D. 1986. 'A Nice Derangement of Epitaphs'. In *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, ed. E. Lepore. Cambridge: Blackwell.
- Davies, M. 1987. 'Tacit Knowledge and Semantic Theory: Can a Five Per Cent Difference Matter?' *Mind*, Vol. 96, No. 384, pp. 441–462.
- Dawkins, R. 1976. *Selfish Gene*. Oxford University Press.

- Dever, J. 2012. 'Formal Semantics'. In García-Carpintero, M., and Kölbel, M. (eds), *The Continuum Companion to the Philosophy of Language*. London: Continuum Ltd.
- Devitt, M. 2006. *Ignorance of Language*. Oxford University Press.
- Devitt, M. 2008a. 'Explanation and Reality in Linguistics'. *Croatian Journal of Philosophy* Vol. VIII, No. 23, 2008.
- Devitt, M. 2008b. 'A Response to Collins' Note on Conventions And Unvoiced Syntax'. *Croatian Journal of Philosophy* Vol. VIII, No. 23, 2008.
- Devitt, M. 2013. 'The "Linguistic Conception" of Grammars'. *Filozofia Nauki* Rok XXI, No. 2(82).
- Devitt, M., and Sterelny, K. 1989. 'Linguistics: What's Wrong with "The Right View"'. *Philosophical Perspectives*, Vol. 3. Philosophy of Mind and Action Theory, pp 497-531.
- Diez, J. 1997a. 'A Hundred Years of Numbers. An Historical Introduction to Measurement Theory 1887-1990—Part 1'. *Studies in History and Philosophy of Science*, 28(1): 167-185.
- Diez, J. 1997b. 'A Hundred Years of Numbers. An Historical Introduction to Measurement Theory 1887-1990—Part 2'. *Studies in History and Philosophy of Science*, 28(2): 237-265.
- Elgin, M., and Sober, E. 2002. 'Cartwright On Explanation And Idealization'. *Erkenntnis*, Vol. 57, No. 3, pp 441-450.
- Epstein, S., and Hornstein, N. 2005. 'Letter on 'The future of language' '. *Language*, 81:3-6.
- Evans, G. 1981. 'Semantic theory and Tacit Knowledge'. in *Collected Papers* 1996, pp 322-342. Clarendon Press.
- Field, H. 1980. *Science without Numbers: A Defense of Nominalism*. Princeton, N.J.: Princeton University Press.
- Field, H. 1985. 'On conservativeness and mathematics'. *The Journal of Philosophy* 82: 239-260.
- Firth, J. 1957. *Papers in Linguistics*. London, Oxford University Press.
- Fodor, J., Bever, T., and Garrett, M. 1974. *The Psychology of Language*. McGraw-Hill Publishers.
- Fodor, J. 1981. 'Some notes on what linguistics is about'. In Block, N. (ed), *Readings in the Philosophy of Psychology* vol. II, Cambridge, Mass: Harvard University Press, pp 197-207.
- Fodor, J. 1983. *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Frigg, R. 2010b. 'Models and Fiction'. *Synthese*, 172(2): 251-268.
- Gamut. 1991. *Logic, Language and Meaning, Volume II: Intensional Logic and Logical Grammar*. University of Chicago Press. Translation and revision of Logica, taal en betekenis II.
- George, A. 1989. 'How not to become confused about linguistics'. In *Reflections on Chomsky* George, A. (ed.), pp 90-110. Basil Blackwell.
- George, A. 1996. 'Katz Astray'. *Mind and Language*, Vol. 11. No. 3, pp 295-305.
- Giere, R. 1988. *Explaining Science: A Cognitive Approach*. Chicago: Chicago University Press.
- Gilbert, M. 1989. *Social facts*. New York: Routledge.
- Glanzberg, M. 2014. 'Explanation and Partiality in Semantic Theory'. In *Metasemantics: New Essays on the Foundations of Meaning* (eds) Burgess, A and Sherman, B. Oxford University of Press.
- Gödel, K. 1944. 'Russell's Mathematical Logic'. In Benacerraf, P., and Putnam, H. (1983) (eds.), *Philosophy of Mathematics: Selected Readings*. Cambridge University Press.

- Godfrey-Smith, P. 2006. 'The Strategy of Model-based Science'. *Biology and Philosophy*, 21: 725–740.
- Hale, R. 1987. *Abstract Objects*. Oxford: Basil Blackwell.
- Hammond, M. 1997. *Parsing syllables: Modeling OT computationally*. Tucson, AZ: University of Arizona.
- Harris, Z. 1951. *Methods in Structural Linguistics*. Chicago: University of Chicago Press.
- Hawthorne, J., and Lepore, E. 2011. 'On Words'. *The Journal of Philosophy*, Vol 108(9): 447–485.
- Heim, I. and Kratzer, A. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.
- Hellman, G. 1989. *Mathematics without Numbers: Towards a Modal-Structural Interpretation*. Oxford: Clarendon Press.
- Heyting, A. 1974. 'Intuitionistic Views on the Nature of Mathematics'. *Synthese*, Vol. 27, No. 1/2, pp. 79–91.
- Higginbotham, J. 1985. 'On Semantics'. *Linguistic Inquiry* 16: 547–93.
- Higginbotham, J. 1989. 'Knowledge of Reference'. In *Reflections on Chomsky*, George, A. (ed.), pp 153–174. Basil Blackwell.
- Higginbotham, J. 1991. 'Remarks on the Metaphysics of Linguistics'. *Linguistics and Philosophy*, Vol. 14, No.5 pp 555–566.
- Hilbert, D. [1899] 2004. *Grundlagen der Geometrie*. Leipzig, Teubner; *Foundations of Geometry*, tr. by Townsend, E, La Salle, Illinois, Open Court, 1959.
- Hintikka, K. 1999. 'The Emperor's New Intuitions'. *The Journal of Philosophy*, Vol 96. No.3: 127–147.
- Hinzen, W. 2000. 'Minimalism'. In R. Kempson, T. Fernando, & N. Asher (Eds.), *Handbook of Philosophy of Science 14: Philosophy and Linguistics*, pp. 93–142. Amsterdam: Elsevier.
- Hinzen, W. 2012. 'Minimalism'. In Kempson, R., Fernando, T., and Asher, N. (eds). *Philosophy of Linguistics*. Elsevier B.V. North Holland.
- Hinzen, W., and Uriagereka, J. 2006. 'On the Metaphysics of Linguistics'. In *Erkenntnis*, Vol. 65, No.1, pp 71–96.
- Hockett, C. 1968. *The State of the Art*. The Hague & Paris: Mouton.
- Hodges, W. 2009. 'Functional Modelling and Mathematical Models: A Semantic Analysis'. In *Philosophy of Technology and Engineering Sciences*. Elsevier BV.
- Hoeksema, J. 1997. *Corpus study of negative polarity items*. University of Groningen. IV-V Jornades de corpus linguistics 1996–1997, Universitat Pompeu Fabre, Barcelona.
- Hughes, R. 1997. 'Models and Representation'. *Philosophy of Science* 64: S325–S336.
- Ishihara, R. 2015. 'Intuitionism in the Philosophy of Mathematics', *The Stanford Encyclopedia of Philosophy*, Zalta, E. (ed.), URL = <http://Plato.stanford.edu/archives/spr2015/entries/intuitionism/>.
- Jackendoff, R. 1977. *X Syntax: a Study of Phrase Structure*. Cambridge, Mass: MIT Press.
- Jackendoff, R. 2002. *The Foundations of Language*. Oxford University Press.
- Jackendoff, R. 2007. 'Linguistics in Cognitive Science: The state of the art'. *The Linguistic Review* 24, pp 347–401.
- Jackendoff, R., and Pinker, S. 2005. 'The nature of the language faculty and its implications for evolution of language - (Reply to Fitch, Hauser, and Chomsky)'. *Cognition* 97: 211–225.
- Jacobson, S. 1984. *Central Yupi'k and the Schools: A Handbook for Teachers*. Juneau, Alaska: Alaska Native Language Center and Alaska Department of Education. alaskool.org

- Jacobson, P. 2014. *Compositional Semantics: An Introduction to the Syntax/Semantics Interface*. Oxford University Press.
- Jäger, G., and Rogers, J. (2012). 'Formal language theory: refining the Chomsky hierarchy.' *Philosophical Transactions of the Royal Society B: Biological Sciences* 367. 1598: 1956–1970.
- Johnson, K. 2007. 'The Legacy of Methodological Dualism'. *Mind and Language*, Vol. 22 No. 4, pp. 366–401.
- Johnson, K. 2015. 'Notational Variants and Invariance in Linguistics'. *Mind and Language*, Vol. 30, No. 2, pp 162–186.
- Kaplan, D. 1990. 'Words'. *Aristotelian Society Supplementary Volume* LXIV: 93–119.
- Katz, J., and Fodor, J. 1963. 'The Structure of Semantic Theory'. *Language*, 39:170–210.
- Katz, J. 1966. *The Philosophy of Language*. New York: Harper and Row.
- Katz, J. 1971. *Linguistic Philosophy: the underlying reality of language and its philosophical import*. George Allen and Unwin Ltd: London.
- Katz, J. 1971. *Linguistic Philosophy: the underlying reality of language and its philosophical import*. George Allen and Unwin Ltd: London.
- Katz, J. 1972. *Semantic Theory*. Harper and Row, New York.
- Katz, J. 1981. *Language and Other Abstract Objects*. Rowman and Littlefield, Totowa, New Jersey.
- Katz, J. 1985. 'An outline of Platonist grammar'. In Katz, J (ed) *The Philosophy of Linguistics*. Oxford University Press.
- Katz, J. 1990. *The Metaphysics of Meaning*. MIT Press, Cambridge, Massachusetts.
- Katz, J. and Postal, P. 1991. 'Realism vs. Conceptualism in Linguistics'. *Linguistics and Philosophy* Vol. 14, No.5 pp 515–554.
- Katz, J. 1995. 'What Mathematical Knowledge Could be'. *Mind*. Vol. 104, No. 415, pp 491–522.
- Katz, J. 1996. 'The Unfinished Chomskyan Revolution'. *Mind and Language*, 11:3, pp 270–294.
- Katz, J. 1998. *Realistic Rationalism*. MIT Press.
- Kaufman, D. 2002. 'Composite Objects and the Abstract/Concrete Distinction'. *Journal of Philosophical Research*, Vol. 27, pp 215–238.
- Keller, F. 1998. 'Grammaticality Judgments and Linguistic Methodology'. Research Paper EUCCS-RP-1998-3.
- Kempson, R., Meyer-Viol, W., Gabbay, D., 2001. *Dynamic syntax - the flow of language understanding*. Blackwell Publishers.
- Kincaid, H. (2008). 'Structural Realism and the Social Sciences'. *Philosophy of Science*, Vol. 75, No. 5, pp 720–731.
- Kirby, S. 1999. 'Learning, Bottlenecks and Infinity: a working model of evolution of syntactic communication'. In Dautenhahn, K. and Nehaniv, C. (eds.) *Proceedings of the AISB'99 Symposium on Imitation in Animals and Artifacts*.
- Klein, U. 1999. 'Techniques of modelling and paper-tools in classical chemistry'. In Morgan, M., and Morrison, M. (eds) *Models as Mediators*, pp 146–167. Cambridge University Press.
- Kripke, S. 1982. *Wittgenstein on Rules and Private Language*. Harvard University Press.
- Ladyman, J. (1998). 'What is Structural Realism?' *Studies in the History and Philosophy of*

- Science*, Vol. 29, No. 3, pp 409-424.
- Lambek, J. 1958. 'The mathematics of sentence structure'. *American Mathematical Monthly* 65:154-170.
- Langendoen, T. 1998. 'Linguistic Theory.' In Bechtel, W., and Graham, G. (eds) *A Companion to Cognitive Science*, MA: Blackwell Publishers, pp 235-244.
- Langendoen, T. 2010. 'Just how big are natural languages?' In van der Hulst, H. (ed.), *Recursion and Human Language*, pp 139-47. De Gruyter Mouton.
- Langendoen, T., and Postal, P. 1984. *The Vastness of Natural Languages*. Blackwell Publishers.
- Lappin, S., Levine, R., and Johnson, D. 2000b. 'The Structure of Unscientific Revolutions'. *Natural language and Linguistic Theory* 18: 665-671.
- Larson, R. and Segal, G. 1995. *Knowledge of Meaning*. Cambridge, MA: MIT Press.
- Lasnik, H. 2000. 'Syntactic Structures'. In *Revisited: Contemporary Lectures on Classic Transformational Theory*. Cambridge, MA: MIT Press.
- Lasnik, H. 2005. 'Grammar, levels, and biology'. In: McGilvray, J. (ed.), *The Cambridge Companion to Chomsky*. Cambridge University Press, Cambridge, pp. 60-83.
- Laurence, S. 2003. 'Is Linguistics a Branch of Psychology?' In *Barber*, 69-106.
- Lenci, A. 2008. 'Distributional semantics in linguistic and cognitive research. A foreword'. *Rivista di Linguistica* 20:1, pp. 1-30.
- Lepore, E., & Ludwig, K. 2007. 'The Reality of Language: on the Davidson/Dummett exchange'. In R. E. Auxier & L. E. Hahn (Eds.), *The philosophy of Michael Dummett* (pp. 185-214). LaSalle: Open Cour.
- Lesmo, L, and Robaldo, L. 2006. 'Dependency tree semantics.' In *Foundations of Intelligent Systems*. Springer Berlin Heidelberg, pp 550-559.
- Lewis, D. 1969. *Conventions: A Philosophical Study*. Harvard University Press.
- Lewis, D. 1970. 'General Semantics'. *Synthese* 22:18-67.
- Lewis, D. 1975. 'Languages and Language'. In Matinich, P. (ed) *The Philosophy of Language*. Oxford University Press.
- Lewis, D. 1980. 'Index, Context, and Content'. In Kanger, S., and Öhman (eds.) *Philosophy and Grammar*, 79-100.
- Linsky, B., and Zalta, E. 1995. 'Naturalized Platonism vs. Platonized Naturalism'. *Journal of Philosophy* xcii/10: 525-555.
- Lobina, D. 2014. 'When linguists talk mathematical logic'. *Frontiers in Psychology*, Vol. 5. 382, pp 1-3.
- Ludlow, P. 2009. 'Review: *Ignorance of Language*'. *The Philosophical Review*, Vol. 118, No. 3, pp. 393-402.
- Ludlow, P. 2011. *The Philosophy of Generative Grammar*. Oxford University Press.
- Mäki, U. 2011. 'Models and the Locus of their Truth.' In *Synthese* 180, pp 47-63.
- Manning, C. 2003. 'Probabilistic Syntax'. In Bod, R., Hay, J., and Jannedy, S. (eds), *Probabilistic Linguistics*. The MIT Press, Cambridge Massachusetts. pp 289-342.
- Marantz, A. 2005. 'Generative linguistics within the cognitive neuroscience of language'. *The Linguistic Review* 22, pp 429-446.
- Marr, D. 1982. *Vision*. W.H. Freeman and Company: New York.
- Matthews, R. 2007. *The Measure of Mind: Propositional Attitudes and Their Attribution*.

- Oxford: Oxford University Press.
- McCarthy, J. (ed.) 2003b. *Optimality Theory in Phonology: A Reader*. Malden, MA, and Oxford, UK: Blackwell.
- McDonald, F. 2009. 'Linguistics, Psychology, and the Ontology of Language'. *Croatian Journal of Philosophy* Vol. IX, No. 27.
- Michaelis, J. 2001. 'Transforming linear context free rewriting systems into minimalist grammars.' *Lecture Notes in Artificial Intelligence*, No. 2099. In: de Groote, P., Morill, G., Restore, C. (Eds.), *Logical Aspects of Computational Linguistics*, pp. 228–244.
- Millikan, R. 2005. *Language: A Biological Model*. Clarendon Press: Oxford.
- Millikan, R. 2005. 'Meaning, Meaning, and Meaning'. In *Language: a Biological model*. Oxford University Press. pp 53–76.
- Mönnich, U. 2007. 'Minimalist syntax, multiple regular tree grammars, and direction preserving tree transducers.' In: Rogers, J., Kepser, S. (Eds.), *Model Theoretic Syntax at 10. ESSLLI'07 Workshop Proceedings*.
- Montague, R. 1970a. 'English as a Formal Language'. In: Bruno Visentini (ed.): *Linguaggi nella società e nella tecnica*. Reprinted in R. Montague, *Formal Philosophy*. (ed.) R. Thomason. New Haven, CT: Yale University Press, 1974, pp 189–223.
- Montague, R. 1970b. 'Universal Grammar'. *Theoria* 36:373–98. Reprinted in R. Montague, *Formal Philosophy*. (ed.) R. Thomason. New Haven, CT: Yale University Press, 1974, pp. 222–46.
- Morgan, M., and Morrison, M. 1999. *Models as Mediators*. Cambridge University Press.
- Morrison, M. 2015. *Reconstructing Reality: Models, Mathematics, and Simulations*. Oxford University Press.
- Nefdt, R. 2016. 'Linguistic Modelling and the Scientific Enterprise'. *Language Sciences*. Vol. 54, pp 43–57.
- Newmeyer, F. 1996. *Generative Linguistics: A Historical Perspective*. London: Routledge.
- Parsons, C. 1980. 'Mathematical Intuition'. *Proceedings of the Aristotelian Society*, New Series, Vol. 80, pp. 145–168.
- Parsons, C. 1990. 'The Structuralist view of mathematical objects'. *Synthese* 84: 303–346.
- Parsons, C. 2004. 'Structuralism and Metaphysics'. *The Philosophical Quarterly*, Vol. 54, No. 214, pp 56–77.
- Partee, B. 1979. 'Semantics – Mathematics or Psychology?'. In Bauerle, U., and von Stechow, A. (eds.) *Semantics from Different Points of View*. Berlin: Springer, pp 1–14.
- Peacocke, 1986. 'Explanation in Computational Psychology: Language, Perception and Level 1.5'. *Mind and Language*, Vol. 1, No. 2, pp 101–123.
- Peacocke, 1989. 'When is a grammar psychologically real'. In *Reflections on Chomsky* George, A. (ed.), pp 111–130. Basil Blackwell.
- Peregin, J. 2015. *Inferentialism: Why Rules Matter*. Palgrave Macmillan.
- Pesetsky, D. 2009. 'Linguistic Universals and Universal Grammar'. The MIT Press.
- Pierce, C. 1958. *Collected Papers of Charles Sanders Peirce*, Vol. 4. Cambridge, MA: Harvard University Press.
- Pincock, C. 2004a. 'A Revealing Flaw in Colyvan's Indispensability Argument'. *Philosophy of Science* 71: 61–79.
- Pincock, C. 2007. 'A Role for Mathematics in the Physical Sciences'. *Nous* 41: 253–275.

- Pinker, S. 1994. *The Language Instinct*. Harper Perennial Modern Classics.
- Portides, D. 2013. 'Idealization in Economic Modeling'. In Andersen, H., Dieks, D., Gonzalez, W.J., Uebel, T., Wheeler, G. (Eds.) *New Challenges to the Philosophy of Science*. Springer Press.
- Postal, P. 2003. 'Remarks on the foundations of linguistics'. *The Philosophical Forum*. Vol. XXXIV, Nos. 3 & 4.
- Postal, P. 2003. *Skeptical Linguistic Essays*. Oxford University Press.
- Postal, P. 2009. 'The Incoherence of Chomsky's 'Biolinguistic' Ontology'. *Biolinguistics* 3.1: 104-123.
- Prince, A., and Smolensky, P. 1993/2004. *Optimality theory: Constraint interaction in generative grammar*. Rutgers University and University of Colorado at Boulder. Technical Report RuCCSTR-2, available as ROA 537-0802. Revised version published by Blackwell, 2004.
- Pullum, G. 1983. 'How Many Possible Human Languages Are There?'. *Linguistic Inquiry*, Vol. 14, No. 3, pp. 447-467.
- Pullum, G. 2011. 'On the mathematical foundations of syntactic structures'. *Journal of logic, language and information*, 20(3), pp 277-296.
- Pullum, G. 2013. 'The Central Question in Comparative Syntactic Metatheory'. *Mind & Language*, Vol. 28, No. 4, pp. 492-521.
- Pullum, G. and Scholz, C. 2001. 'On the distinction between model-theoretic and generative-enumerative syntactic framework.' In *Logical Aspects of Computational Linguistics: 4th International Conference*. Springer Press.
- Pullum, G. and Scholz, B. 2010. 'Recursion and the infinitude claim'. In van der Hulst, H. (ed.), *Recursion in Human Language (Studies in Generative Grammar 104)*, pp 113-138. Berlin: Mouton de Gruyter.
- Pustejovsky, J. 1998. *The generative lexicon*. The MIT Press, Cambridge & London.
- Pylyshyn, Z. 1973. 'What the mind's eye tells the mind's brain: a critique of mental imagery.' *Psychological Bulletin*, 80: 1-24.
- Pylyshyn, Z. 1984. *Computation and Cognition*. Cambridge, MA: MIT Press.
- Pylyshyn, Z. 1991. 'Rules and representations: Chomsky and representational realism.' In Kasher, A. (ed.), *The Chomskyan turn*. Cambridge, MA : Blackwell, pp 231-51.
- Quine, W.V. 1972. 'Methodological reflections on current linguistic theory'. In Davidson, D., and Harman, G. (eds), *Semantics of Natural Language*. Dordrecht: D. Reidel, pp 442-454.
- Quine, W.V. 1976. *Ways of Paradox and Other Essays*. Harvard University Press.
- Quine, W.V. 1987. *Quiddities*. Cambridge, MA: Harvard University Press.
- Rambow, O., Joshi, A. 1997. 'A formal look at dependency grammars and phrase structure grammars, with special consideration of word-order phenomena'. In Wanner, L (ed.), *Recent Trends in Meaning-Text Theory*, p. 167-190. Amsterdam and Philadelphia: John Benjamins.
- Rayo, A. 2007. 'Ontological Commitment'. *Philosophy Compass* 2/3, 428-444.
- Rayo, A. 2013. 'A Plea for Semantic Localism'. *Nous*, 47:4, 647-679.
- Resnik, M. 1981. 'Mathematics as a Science of Patterns: Epistemology'. *Nous*, Vol.16 pp 95-105.
- Resnik, M. 1981. 'Mathematics as Science of Patterns: Ontology and Reference'. *Nous*, Vol. 15, pp 529-550.
- Resnik, M. 1997. *Mathematics as a Science of Patterns*. Clarendon Press: Oxford.

- Rey G. 2006b. 'Conventions, Intuitions and Linguistic Inexistents: A Reply to Devitt', *Croatian Journal of Philosophy*, 6, pp. 549–569.
- Rogers J. 1998b. 'A Descriptive Characterization of Tree-Adjoining Languages'. *Proc. of the 17th International Conference on Computational Linguistics (COLING'98) and the 36th Annual Meeting of the Association for Computational Linguistics (ACL'98)*, ACL, Montreal.
- Rothschild, D. 2009. 'Definite Descriptions and Negative Polarity'. MS.
- Sampson, G. 2001. *Empirical Linguistics*. Continuum Press.
- Sag, I., Wasow, T., and Bender, E. 2003. *Syntactic Theory: A Formal Introduction. Second Edition*. CSLI Publications.
- Savitch, W. 1993. 'Why it might pay to assume that languages are infinite'. In *Annals of Mathematics and Artificial Intelligence* 8, 17–25.
- Searle, J. 1974. 'Chomsky's Revolution in Linguistics'. In Harman, G. (ed) *On Noam Chomsky: Critical Essays*, pp 2–33.
- Shapiro, S. 1983. 'Conservativeness and Incompleteness'. *The Journal of Philosophy*, Vol. 80, No. 9, pp. 521–531.
- Shapiro, S. 1997. *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press.
- Shapiro, S. 2005. 'Categories, Structures, and the Frege–Hilbert Controversy: The Status of Metamathematics'. *Philosophia Mathematica* III, 13, pp 61–77.
- Shieber, S. 1985. 'Evidence against the context-freeness of natural language'. In *The Formal complexity of natural language*. Springer Netherlands, 1987. 320–334.
- Shieber, S., and Schabes, Y. 1991. 'Synchronous tree-adjoining grammars'. In *Proceedings of the 13th International Conference on Computational Linguistics*, Helsinki.
- Slezak, P. 2007. 'Linguistic Explanation and "Psychological Reality"'. Manuscript online.
- Smith, N., and Tsimpli, I. 1995. *The Mind of a Savant: Language, Learning and Modularity*. Wiley-Blackwell.
- Smith, N., Tsimpli, I., Morgan, G. Woll, B. 2010 *The Signs of a Savant Language Against the Odds*. Cambridge University Press.
- Smolensky, P. 2001. 'Optimality Theory: Frequently Asked "Questions"'. In Fukazawa, H., and Kitahara, M. (eds), *Gengo*. Tokyo: Taishukan.
- Soames, S. 1984a. "What is a Theory of Truth?" *Journal of Philosophy* 81: 411–29.
- Soames, S. 1984b. 'Linguistics and Psychology'. *Linguistics and Philosophy* 7, 155–179.
- Stabler, E. 1997. 'Derivational Minimalism'. In Retoré, C. (ed.), *Logical Aspects of Computational Linguistics*, 68–95. Springer.
- Stabler, E. 2011. 'Meta-meta-linguistics'. *Theoretical Linguistics*. Vol. 37, no. 1/2: 69–78.
- Stainton, R. 1996. *Philosophical Perspectives on Language*. Peterborough, ON: Broadview Press.
- Stainton, R. 2006. 'Meaning and Reference: Some Chomskyan Themes'. In E. Lepore and B. Smith (eds.) *Handbook of the Philosophy of Language*. Oxford: Oxford University Press, pp. 913–940.
- Stainton, R. 2011. 'The Role of Psychology in the Philosophy of Language'. In D. Graff Fara & G. Russell (eds.) *Routledge Companion to the Philosophy of Language*. London: Routledge, pp. 525–532.
- Stainton, R. 2012. 'In Defense of Public Languages'. *Linguistics and Philosophy* 34(5):

479–488.

Stainton, R. 2014. 'Philosophy of Linguistics'. *Oxford Handbooks Online*.

Stanley, J. and Szabó, Z. 2000. 'On Quantifier Domain Restriction'. *Mind & Language*, Vol. 15, No. 2-3, pp 219–261.

Steiner, M. 1998. *The Applicability of Mathematics as a Philosophical Problem*. Cambridge, MA: Harvard University Press.

Strevens, M. 2007. 'Why Explanations Lie: Idealization in Explanation.' Unpublished manuscript, Department of Philosophy, New York University.

Suárez, M. 2004. 'An Inferential Conception of Scientific Representation'. *Philosophy of Science* 71: 767–779.

Suárez, M., and Cartwright, N. 2008. 'Theories: tools versus models'. *Stud. Hist. Philos. Sci. Part B Stud. Hist. Philos. Mod. Phys.* 39 (1), 62–81.

Suppe, F. 1989. *The Semantic Conception of Theories and Scientific Realism*. Urbana: University of Illinois Press.

Suppes, P. 1960. 'A Comparison of the Meaning and use of Models in Mathematics and the Empirical Sciences'. *Synthese*, 12, pp. 287–300.

Szabó, Z. 1999. 'Expressions and Their Representations'. *The Philosophical Quarterly*, Vol. 49, No. 195, pg 145–163.

Szabó, Z. 2012. 'The case for compositionality'. In Werning, M., Hinzen, W., and Machery, E. *The Oxford Handbook of Compositionality*, pp 47–63. Oxford University Press.

Szabó, Z. 2015a. 'Major Parts of Speech'. *Erkenntnis* Vol. 80: 3–29.

Szabó, Z. 2015b. 'Semantic Explanations'. Online draft.

Thatcher, J., and Wright, J. 1968. 'Generalized finite automata theory with an application to a decision problem of second order logic'. *Mathematical Systems Theory*, 2(1):57–81.

Thomason, R. 1974: *Formal Philosophy. Selected Papers by Richard Montague*. New Haven.

Thomson-Jones, M. 2005. 'Idealization and Abstraction: A Framework'. In Cartwright, N., and Jones, M. (eds) *Correcting the Model: Idealization and Abstraction in the Sciences*, pp 173–218. Rodopi Press.

Thomson-Jones, M. 2012. 'Modelling without Mathematics'. *Philosophy of Science*, Vol. 79, pp. 761–772.

Tiede, H., and Stout, L. 2010. 'Recursion, Infinity and Modeling'. In *Recursion and Human Language*. (ed) van der Hulst, H. Berlin/New York: Mouton de Gruyter, pp. 147–158.

Tomalin, M. 2006. *Linguistics and the Formal Sciences*. Cambridge University Press.

Tomalin, M. 2007. 'Reconsidering Recursion in Syntactic Theory'. *Lingua* 117, pp 1784–1800.

van den Bogaard, A. 1999. 'Past measurements and future prediction'. In Morgan, M., and Morrison, M. (eds) *Models as Mediators*, pp 282–325. Cambridge University Press.

van Fraassen, B. 1980. *The Scientific Image*. Oxford: Oxford University Press.

Weisberg, M. 2007a. 'Who is a Modeler?' *British Journal of the Philosophy of Science*, Vol. 58: 207–233.

Weisberg, M. 2007b. 'Three Kinds of Idealization'. *The Journal of Philosophy*, Vol. 104, No. 12: 639–659.

Weisberg, M. 2013. *Simulation and Similarity: Using Models to Understand the World..*

Oxford University Press.

Wetzel, L. 2009. *Types and Tokens*. Cambridge: MIT Press.

Wetzel, L. 2014. 'Types and Tokens', *The Stanford Encyclopedia of Philosophy*, Zalta, E. (ed.), URL = <<http://Plato.stanford.edu/archives/spr2014/entries/types-tokens/>>.

Wigner, E.P. 1960. "The Unreasonable Effectiveness of Mathematics in the Natural Sciences". *Communications on Pure and Applied Mathematics* 13:1–14.

Wittgenstein, L. 1983 [1956]. *Remarks on the Foundations of Mathematics*. The MIT Press.

Williamson, T. 'Model-Building in Philosophy'. In *Philosophy's Future: The Problem of Philosophical Progress*. (forthcoming). Oxford: Wiley.

Wright, C. 1989. 'Wittgenstein's Rule-following Considerations and the Central Project of Theoretical Linguistics'. In *Reflections on Chomsky* George, A (ed.), pp 231–264. Basil Blackwell.

Yablo, S. 2013. 'Explanation, Extrapolation, and Existence'. *Mind*, 121(484), 1007–1029.

Yalcin, S. 2014. 'Semantics and Metasemantics in the Context of Generative Grammar'. In *Metasemantics: New Essays on the Foundations of Meaning* (eds) Burgess, A and Sherman, B. Oxford University of Press.

Yang, C. 2006. *The Infinite Gift*. New York: Scribner.

Zeijlstra, H. 2004. *Sentential negation and Negative Concord*. Ph.D. dissertation, University of Amsterdam.