Structural Realism and Generative Linguistics

Ryan M. Nefdt*

Department of Philosophy

University of Cape Town

November 16, 2020

Abstract

Linguistics as a science has rapidly changed during the course of a relatively short period. The mathematical foundations of the science, however, present a different story below the surface. In this paper, I argue that due to the former, the seismic shifts in theory over the past 80 years opens linguistics up to the problem of pessimistic meta-induction or radical theory change. I further argue that, due to the latter, one current solution to this problem in the philosophy of science, namely structural realism (Ladyman 1998, French 2006), should be viewed as especially enticing for linguists, as their field is a largely structural enterprise. I discuss particular historical instances of theory change in generative syntax before investigating two views on the nature of structural properties and eventually proposing an approach in terms of invariance (Johnson 2015) as a grounding for structural realism in the history and philosophy of linguistics.

Keywords: *philosophy of linguistics, structural realism, generative grammar, syntax, structural properties*

1 Introduction

The generative study of natural language was established in the late 1950's around the distinction between linguistic competence and performance, the former amenable to precise mathematical investigation, while the latter perhaps only to statistical approximation. Since then, generative linguistics has enjoyed much success along a path chartered

*Final draft for article accepted for publication in *Synthese* 2020

with countless discoveries from the formal sciences as applied to the modelling of natural language. At the centre of the newly established discipline was the syntactic engine, the structures of which were revealed through modelling grammatical form. The generativist paradigm in linguistics initially relied heavily upon the proof-theoretic techniques introduced by Emil Post and other logicians to model the form language takes (Tomalin 2006, Pullum 2011, Pullum 2013).[1] Yet despite these aforementioned formal beginnings (and successes), the theory of linguistics has changed its commitments quite drastically over the years, eschewing among other things formalisation, cognitive science for evolutionary biology, derivations for constraints, rules for schemata, phrases for phases and other theoretical moves.

Given significant theory change, the fruitfulness of the enterprise and its erstwhile discoveries are inevitably called into question (Stokhof and van Lambalgen 2011, Lappin et al. 2000, Jackendoff 2002).[2] Thus, the goal of the paper is to argue that adopting the structural realist framework for linguistics addresses this and other philosophical problems. Not only can the view explain radical theory change but it can also offer some resolution to the conflict over the ontology of natural languages in a way consistent with accounts of the natural sciences.

Thus, in this paper, I argue that linguistics as a science essentially faces the problem of pessimistic meta-induction, albeit at a much faster rate than the more established sciences such as physics. In addition, I claim that the focus on the ontology of linguistic objects, such as words, phrases, sentences etc. belies the formal nature of the field which is at base a structural undertaking. Both of these claims, I argue, lead to the interpretation of linguistics in terms of *ontic* structural realism in the philosophy of science (Ladyman 1998, French 2006). Thus, to be realist in this sense is to accept the existence of linguistic structures (not individual objects) defined internally through the operations of the grammars (or another means to be discussed later) and what remains relatively stable across various theoretical shifts in the generative paradigm, from Standard Theory (1957-1980)

[1] Here my focus will largely be on the formal history of generative syntax. A broader view could take the present methodology and extend it beyond generative grammar to the so-called ‘structuralist’ movement of Harris, Bloomfield, and Hockett. See Joseph (1999) and Matthews (2001) for the direct connections between this latter paradigm and contemporary linguistics. See also Nefdt (2019b) for a related account.

[2] A related, more ontological, question is if the grammars of linguistics are scientific theories (as Chomsky and others have insisted over the years), then what are the objects being explained by these grammars? The radical theory change question has received very little attention, while this latter question has received perhaps too much. For instance, Chomsky (1986a) details the received or psychological take on the ontology of linguistics, Katz and Postal (1991) offers a Platonist interpretation, Devitt (2006) a non-psychological physicalist view, and Stainton (2014) a mixture of all the above.

to the Minimalist Program (1995-present), are the structures so defined.

The paper is separated into three parts. In the first part, I focus on some important theoretical shifts which the generative linguistic tradition has undergone since its inception in the late 1950's. For instance, the move from rewriting systems with transformations to X-bar representation (Chomsky 1970) with theta roles to the current single movement operator Merge contained only by constraints. Despite appearances, I hope to show that the general structure of these representations have remained relatively constant. In the second part, I discuss both realism and structural realism in the philosophy of science more generally and why the latter might serve as an illuminating foundation for linguistics, assuming the former. Linguistics here is interpreted structurally without recourse to the independent existence of individual objects in that structure (along the lines of Shapiro 1997 for mathematics). In other words, there are no phrases, clauses or sentences outside of the overarching linguistic structure described by the grammar. Lastly, I delve into the issue of structural properties, detail two distinct approaches to their characterisation, namely definitional and invariance, and follow Johnson (2015) in suggesting the latter as a useful tool in defining the structures of linguistic analysis.

2 Linguistic Theory Change

The history of science bears witness to a number of radical theory changes from Newtonian physics to Relativistic, from Euclidean geometry to Riemannian as a characterisation of physical space, from phlogiston theory to Lavoiser's oxygen theory, among countless others. In the course of such changes, one might easily dismiss the old theory as simply false. Laudan (1981) famously proposed that there might be a deeper issue at stake here, namely what has become known as pessimistic meta-induction (PMI). PMI can be defined as follows for present purposes.

PMI : If all (most) previous scientific theories have been shown to be false, then what reason do we have to believe in the truth of current theories?

The problem with radical theory change is that it causes serious tension for any realist theory of science, which wants to hold to the truth or approximate truth of current theories. Of course, false theories can be responsible for true ones through some sort of trial-and-error process. But the idea that our best current theories are of mere instrumental value for later truth is hard to accept.[3] Furthermore, at no point will certainty naturally force itself upon us, especially since success is not a guarantee of truth (e.g. classical mechanics is still a useful tool for modelling physical phenomena). PMI has an ontological component as well. When theories do change, they often propose distinct and incompatible entities in their respective ontologies. Consider the move from phlogiston theory to oxygen theory. In fact, the term 'phlogiston' has become synonymous with a theoretical term which does not refer to anything.[4] Essentially, the ontological status of the objects of the theories are rendered problematic when radical theory change occurs, which prompts a challenge again to the realist. '[I]f she can't establish the metaphysical status of the objects at the heart of her ontology, how can she adopt a realist attitude towards them?' (French, 2011: 165).

Linguistics too has seen its fair share of radical shifts in theory and perspective over the past few decades. In fact, the early generative tradition of Chomsky (1957) had a more formal mathematical outlook. Drawing inspiration from the work of Emil Post on canonical production systems which are distinctively proof-theoretic devices in which symbols are manipulated via rules of inference in order to arrive at particular formulas (not wholly unlike natural deduction systems), linguistics approached language from a more syntactic perspective.[5] This was due in part to two assumptions, namely that (1) syntax is autonomous from semantics, phonology etc. and (2) that syntax or the form of language is more amenable, than say semantic meaning, to precise mathematical elucidation. Mathematical models of this sort would be a key tool in early generative linguistic analysis. Chomsky (1957: 5) stated the formal position in the following way at the time.[6]

> Precisely constructed models for linguistic structure can play an important role, both positive and negative, in the process of discovery itself. By push-

[3]There are such instrumentalist theories on the market. See van Fraasen (1980) constructive empiricism as one prominent example. A general problem for such views is that they tend to make miraculous the explanatory and predictive successes of scientific theory. Van Fraasen's response to these sorts of worries is to appeal to an analogy with evolutionary theory such that only the fittest theories survive (where 'fittest' means something like 'latching on to actual regularities in nature') (van Fraasen, 1980: 40).

[4]In section 5, we discuss Ladyman's (2011) account of the structural continuity of the otiose phlogiston theory more closely.

[5]For a thorough discussion of the influence of Post on generative grammar, see Pullum (2011) and Lobina (2017).

[6]I attempt to follow Pullum and Scholz (2007) throughout in slaloming my way through the minefield of the distinctions between 'formalisation', 'formal', and 'Formalism'. The senses expressed here are related to 'formal' as a term used for systems which abstract over meaning and 'formalisation' as a tool for converting statements of theory into precise mathematical representations. Early generative grammar can be seen as a theory which aimed to achieve both distinct goals.

> ing a precise but inadequate formulation to an unacceptable conclusion, we can often expose the exact source of this inadequacy and, consequently, gain a deeper understanding of the linguistic data. More positively, a formalized theory may automatically provide solutions for many problems other than those for which it was explicitly designed.

He goes on to chastise linguists who are skeptical of formal methods. However, as we shall see, the course of linguistic theory saw a decrease in formalisation and an increased resistance to it (partly inspired by Chomsky's later views). In fact, a generative grammar in the early stages was expressly noncommittal on ontological questions. 'Each such grammar is simply a description of a certain set of utterances, namely, those which it generates' (Chomsky, 1957: 48). By the 1960's, grammars were reconceived as tools for revealing linguistic competence or the idealised mental states of language users. With mentalism, linguistics looked towards sciences such as psychology, physics, and biology for methodological guidance as opposed to logic and mathematics as it had before. As Cowie (1999: 167) states of the time after *Aspects* '[Chomsky] seemed also to have found a new methodology for the psychological study of language and created a new job description for linguists'. The psychological interpretation of linguistic theory held sway until the 1990's when the "biolinguistic" program emerged as yet another new way of theorising about language.[7] The *Minimalist Program* (1995) pushed the field towards understanding language as a 'natural object' in which questions of its optimal design and evolution take centre-stage.[8]

Each new foundation distanced itself from the methodology of its predecessor, postulated different objects and advocated different ends. Thus, PMI takes on special significance for linguistics and an answer to the puzzles it presents become especially peremptory in this light. In the following sections, I will focus on some specific cases of the methodological changes which underlie the above picture.

[7] Of course, the term dates back to Lenneberg (1967) who introduced these issues to the generative linguistics community.

[8] Matters are not as simple as suggested here. As Bickerton (2014) stresses, the peculiarity of the situation in linguistics is that the field at present still contains scholars working in various versions of the generative programme concurrently.

3 From Phrase Structure to Phase Structure

In this section, I aim to provide a story of the mathematical formalisms employed in the service of an ever-changing landscape of theory in linguistics. Many of the theoretical postulates, such as 'deep structure vs surface structure', the modules of Government and Binding theory, domain specificity, optimality, I-language etc., of various generations of generative grammar are not explicitly dealt with as such a narrative would entail more space and a less circumscribed purpose than I have in the present work.

The early generative approach had a particular notion of a language and accompanying grammar at its core. On this view, a language L is modelled on a formal language which is a set of strings characterisable in terms of a grammar G or a rule-bound device responsible for generating well-formed formulas (i.e. grammatical expressions). In *LSLT*, Chomsky (1975: 5) writes of a language that it is 'a set (in general infinite) of finite strings of symbols drawn from a finite "alphabet" '. In formal language theory (FLT) (which took inspiration from this period), assuming a start symbol S, set of terminals (words) T, non-terminals NT (syntactic categories) and production rules R, we can define a grammar in the following way.

> 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 $\mathscr{L}(G)$ (Jäger and Rogers, 2012: 1957).

In Chomsky (1956), natural languages were shown to be beyond the scope of languages with production rules such as $A \rightarrow a$, $A \rightarrow aB$ or $A \rightarrow \varepsilon$ (ε is the empty string) such that $A, B \in NT$ and $a \in T$ (i.e. regular languages).[9] This result lead to the advent of phrase-structure or context-free grammars with production rules of the following sort: either $S \rightarrow ab$ or $S \rightarrow aSb$ (read the arrow as 'replace with' or rewrite). These grammars can handle recursive structures and contain the regular languages as a proper subset. For many years, phrase-structure grammars were the standard way of describing linguistic phenomena. Essentially, phrase structure grammars are rewriting systems in which symbols are replaced with others such as $S \rightarrow NP, VP$ or $NP \rightarrow det, N'$. As Freidin notes 'phrase structure rules are based on a top-down analysis where a sentence is divided into its major con-

[9] Basically, regular grammars can't handle constructions like centre embeddings such as *The boy the girl loved left.* These latter constructions form part of a larger class of non-serial dependencies which are inaccessible to regular languages.

stituent parts and then these parts are further divided into constituents, and so on until we reach lexical items' (2012: 897). There are a number of equivalent means of representing the structure of sentences in this way. The most common is *via* hierarchical diagrams, shown below.

1. 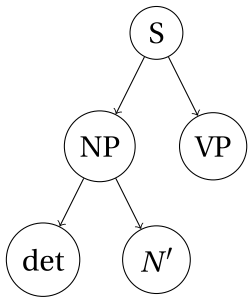

Alternatively one can capture the same information as:

2. $[_S[_{NP}[_{det}][_{N'}]][_{VP}]]$

This basic structure, however, proved inadequate as a means of capturing the structure of passives and certain verbal auxiliary constructions as shown originally in Postal 1964. Transformations were meant to buttress the phrase structure system in order to bridge this gap in explanation. Transformation rules operate on the output of the phrase structure rules and create a derived structure as in (3) below for passivization.

3. NP_1 V NP_2 → NP_2 be-en (AUX) V NP_1

The combined expressive power of phrase structure and transformations proved very productive in characterising myriad linguistic structures. This productivity, with its increased complexity, however, came at a cost to learnability. '[I]f a linguistic theory is to be explanatorily adequate it must not merely describe the facts, but must do so in a way that *explains* how humans are able to learn languages' (Ludlow, 2011: 15). The move to more generality led in part to the Extended Standard Theory and the X-bar schema.

Since the continued proliferation of transformations and phrase structure rules were considered to be cognitively unrealistic, linguistic structures needed more sparse mathematical representation. Although, as Bickerton (2014: 24) states 'rule proliferation and "ordering paradoxes" were only two of a number of problems that led to the eventual replacement of the Standard Theory'.[10]

[10] 'Ordering paradoxes' here refer to the situation in which there are equally valid reasons for orderings from X to Y and Y to X despite the grammar requiring a particular order to pertain.

There was also some theoretical push for more general structure from the Universal Grammar (UG) postulate assumed to be the natural linguistic endowment of every language user. UG needed to contain more general rule schemata in order to account for the diversity of constructions across the world's languages. This structural agenda dovetailed well with the Principles and Parameters (P&P) framework which posited that the architecture of the language faculty constituted a limited number universal principles constrained by individual parametric settings, where 'parameters' were roughly the set of possible variations of a given structure. For instance, some languages such as English require a mandatory NP/DP in the subject position of sentences whereas in pro-drop languages, such as Spanish, empty categories can do the job.

4. It is raining.

5. Lluevé.

These kinds of differences could be expressed in the language of parametric settings. The so-called Extended Projection Principle might be universal but certain languages can contain distinct parameters with relation to it (such as fulfilling it with a null determiner). In other words, a child in the process of acquiring her first language can 'set' the parameter based the available linguistic environment in which she found herself, like flicking a switch. Furthermore, this kind of structural picture is represented well in the X-bar schema (Jackendoff 1977) which contains only three basic rules. There is (1) a specifier, (2) an adjunct, and (3) a complement rule. The idea is that the schema effectively treats endocentric projection as an axiom, which the previous phrase structure rules did not. "Endocentric" here roughly means that one element (i.e. the head) of a constituent determines the function and nature of the whole. The X-bar schema, in other words, restricts the class of phrase markers available (this was part of Chomsky's (1970) original motivation at least).

The specifier rule is given below (where X is a head-variable and XP and YP are arbitrary phrasal categories determined by that head).

- Specifier rule: $XP \rightarrow (Spec)X'$ or $XP \rightarrow X'(YP)$

Or equivalently:

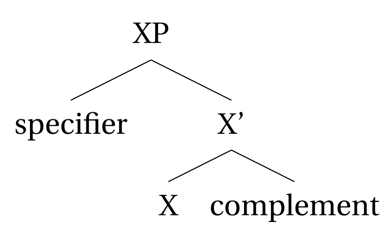

A vast amount of linguistic structure can be modelled by means of this formalism.[11] In fact, X-bar theory over-generates structural descriptions (which need to be reined in by various constraints). But the underlying idea is that our mental competence is more likely to contain generalised rule schemata such as those above than individual phrase structure rules and countless transformations for each natural language. In a sense, X-bar merely smooths over the individual hierarchical structures of before and homes in on a more abstract structural representation for language. As Poole (2002: 50) mentions:

> [W]e discovered that your language faculty appears to structure phrases into three levels: the maximal projections or XP level, the intermediate X' level, and the head or $X°$.

These rules subsume the previous *ad hoc* phrase structure rules. Importantly, the representation, however, only allows for binary rules (unlike the possible n-ary branches of phrase structure trees). Freidin (2012) further claims that X-bar theory represented a shift from top-down to bottom up analysis, despite being formulated in a top down manner a decade into its inception. Here, the idea is that the rules stated above are projections from lexical items to syntactic category labels not the other way around.

Unfortunately, history has a way of repeating itself. Where in the previous instantiation of generative grammar, the proliferation of transformations became unweildy, parameters would soon see a similar fate befall its fecundity. Briefly, UG was assumed to be extremely rich during this period, 'the available devices must be rich and diverse enough to deal with the phenomena exhibited in the possible human languages' (Chomsky, 1986a: 55). However, what was innate and what was learned or set by experience relied in part on a distinction between 'core' grammar and 'periphery', never explicitly provided by the theory (see Pullum (1983) and Culicover (2011) for discussion). Although, formally all previous transformations were reduced to the 'move alpha' operation, the multiplication of parameters took similar shape to its transformational predecessor. Newmeyer (1996: 64) describes this period as one of instability and confusion.

> In the worst-case scenario, an investigation of the properties of hundreds of languages around the world deepens the amount of parametric variation postulated among languages, and the number of possible settings for each param-

[11] I more or less follow the standard story here but see Kornai and Pullum (1990) for a series of convincing arguments to the effect that the X-bar formalism lacks substance in terms of illuminating phrase structure properties without significant restructuring (which they provide).

> eter could grow so large that the term ‘parameter’ could end up being nothing more than jargon for language-particular rule.

What’s more is that these parameters seemed to force the violation of the binary requirement set by the X-bar formalism and with it the cognitive plausibility transiently acquired after the Standard Theory. There needed to be a better way of capturing the movement toward simplifying the grammatical representation and theory of natural language syntax. This and other theoretical motivations led to the Minimalist Program (1995) which pushed the new biolinguistic agenda and a call for further simplicity.

As mentioned in section 2, the question of the evolution of language reset the agenda in theoretical linguistics at this time. The grammatical formalisms assumed to underlie the cognitive aspects of linguistic competence were forced to change with this new perspective, with the result that many of the advances made by the P&P and Government and Binding (1981) theories needed to be abandoned (according to Lappin et al. 2000).[12] Of course, abandonment is a strong claim. Many linguists consider GB to have been on the right track but too complex in its analysis while MP merely filters the structures to only involve the “conceptually necessary” (again, in line with the structural realist interpretation I proffer below).

The rationale was something of the following sort.

> Evolutionarily speaking, it is hard to explain the appearance of highly detailed, highly language-specific mental mechanisms. Conversely, it would be much easier to explain language’s evolution in humans if it were composed of just a few very simple mechanisms (Johnson, 2015: 175).

The Merge operation represented the goal of reducing structure to these simple mechanisms. In the Standard and Extended theories, grammars followed the structures set by

[12] There are some linguists who resist this claim. Instead they claim theoretical continuity between the programmes. For instance, Hornstein (2009) offers two reasons for the theoretical continuity between Minimalism and GB.

> First, MP starts from the assumption that GB is roughly correct. It accepts both the general problems identified for solution (e.g. Plato’s Problem) and the generalizations (“laws”) that have been uncovered (at least to a good first approximation). The second way that MP continues the GB program is in its identification with the Rationalist research strategy that sits at the core of Chomskyan enterprise in general and GB in particular (178).

This might indeed be the case but in my view can best be described as a theoretical orientation rather than theoretical commitment. Many very different theories can be described as “rationalist” in this broad sense. I also worry about the veracity of the first reason but further discussion will take us into exegetical territory.

the proof theory in the early 20th century (see above) which often resulted in grammars 'of roughly the order of complexity of what is to be explained' (Chomsky, 1995: 233). In the Minimalist programme, this apparatus was reduced to a simple set-theoretic operation which takes two syntactic objects and creates a labelled output of their composition (the label to be determined by the features of the objects thereby replacing the projection from heads of X-bar theory).[13] The formulation is given below:

7. $\text{Merge}(\alpha, \beta) = \{\gamma, \{\alpha, \beta\}\}$

Or again, equivalently:

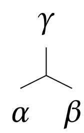

The above is an example of external set merge (where γ is a label projected from one of the elements). Internal merge accounts for recursive structures since it applies to its own output (as in if β is already contained in α). Consider the following sentence.

8. The superhero should fly gracefully.

In a bottom up fashion, *fly* and *gracefully* will merge to form a VP, thereafter this union will merge with the auxiliary *should* to form a TP or Tense Phrase. Merge will independently take *the* and *superhero* and create an NP which will merge to form the final TP to deliver (8) above (the T is the label projected for the entire syntactic object). Importantly for the proposal I will present, '[t]his last step merges two independent phrases in essentially the same way that generalized transformations operated in the earliest transformational grammars' (Freidin, 2012: 911).[14] Thus, although the phrase structure rules had been replaced by the less complex merge operation with *phases*, which are cyclic stages applying to the innermost constituents of the entire process (Chomsky 2008), the structure is identical in the derivation.

[13]Technically, as Langendoen (2003) notes 'Merge is not a single operation, but a family of operations. To belong to the merge family, an operation must be able to yield an infinite set of objects from a finite basis' (307). However, by this definition, the phrase structure rules with recursive components would also be included. The structural similarities of various versions of this infinity requirement on grammars will be discussed in the next section.

[14]The practice of taking ideas or insights in some disguised form from early frameworks is not uncommon. For example, the binding theory of Government and Binding is very close (if not identical) to principles governing anaphora (like the Ross-Langacker constraints) that were first articulated in the 1960's. Similarly, the trace theory of movement is closely tied to the earlier idea of global derivational constraints.

Of course, unlike the top-down analysis of early generative grammar, Merge operates from lexical items in the opposite direction (Merge and the 'lexical array' constituting 'narrow syntax', see Langendoen 2003). As shown in the example above, it does apply to more complex units and their outputs. However, as Lobina (2017) cautions 'talk of top-down and bottom-up derivations is clearly metaphorical' (84).[15] It might add something in appreciating the flavour of the computational process at hand, but often the overall structural picture is unchanged by such parlance.

Lastly, the notion of a *phase* is relevant here. A phase is created when the construction of a constituent XP is followed by access to the lexicon. This can occur when a lexical item can be inserted into a matrix CP (complementizer phrase) in cases in which earlier insertion, in an embedded CP, would have delayed movement. More importantly for our purposes, from the definition of a phase, we get the *Phase Impenetrability Condition* or the claim that if X is dominated by a complement of a phase YP, X cannot move out of YP.

Although phase theory was introduced in Chomsky (1998), one aspect of its structure predates this introduction by three decades, namely so-called 'island effects' (Chomsky 1964, Ross 1967). This is a massive topic in linguistics, so I will briefly focus on the Wh-island constraint and its similar treatment in early generative grammar and by means of phases in the more contemporary setting here. Consider the two sentences below:

9. Which book did Sarah say Mary liked?

10. * Which book did Sarah wonder whether Mary liked?

The above examples show a few things about the structure of Wh-movement. Movement itself is generally taken to be unbounded but there are structures that can block it. For instance, (10) shows that Wh-movement can be blocked in embedded clauses containing *whether*. Both (9) and (10) show that movement happens in small steps (from CP to CP) since if it happened in a big step from the bottom of the tree in (9), then (10) should be licensed likewise.

Island effects were initially explained by means of the A-over-A principle or 'if a rule ambiguously refers to A in a structure of the form of (i), the rule must apply to the higher, more inclusive, node A' (Chomsky 1964).

[15]Compare this metaphorical language to a similar caution in Pullum (2013: 496), '[t]he fact that derivational steps come in a sequence has encouraged the practice of talking about them in procedural terms. Although this is merely a metaphor, it has come to have a firm grip on linguists thinking about syntax'.

i ...$[_{A}...[_{A}...]$

ii 1. I won't read $[_{NP}$ the book on $[_{NP}$ syntax]].
2. *Syntax, I won't read the book on
3. The book on syntax, I won't read

The embedded NP in (ii.1) is blocked from moving in (ii.2) by the principle (later subsumed under the Empty Category Principle or ECP). The island blocks the movement, where an "island" is understood as a constituent that "traps" items from moving out of them.

But this phenomenon can be explained in terms of phases as well.[16] A Wh-island arises when the SpecCP in the middle is already full. Since the Wh-word in the embedded clause cannot be moved into SpecCP, it gets trapped. The CP phase completes, and the higher interrogative C can no longer access the wh-word because it is inside of a finished phrase as in (11).

11. *Which book$_i$ did Sarah think who$_j$ [who$_j$] wanted to read [which book$_i$]?

The explanatory strategy involves certain structural configurations which block the movement of items in embedded units or phrases. Another way of capturing this is that certain phases (CPs or vPs) do not allow Wh-movement to proceed through their specifiers ($Spec$). These phases are then the islands. There is a clear shift from the definition of islands in the A-over-A principle to their definition as phases via the Phase Impenetrability Condition in Minimalism. Despite this, the strategies for dealing with Wh-islands are similar from a structural point of view (as will be argued below).

Let this serve as an account of some of the formal and theoretical changes of generative grammar over the 80 year period since its inception. Below, I will draw on the picture developed here to argue for the structural continuity of linguistics despite the theoretical shifts the overarching theory might have taken during this time.

4 Why Realism?

Before motivating an account which aims to address the PMI while attempting to respect the nature of generative linguistic theorising, a preliminary question needs to be asked

[16] Of course, the immediate predecessor of phases can be found in *barriers*. See Chomsky 1986b for more details on the general framework.

and answered. Why *should* we be realists about linguistics in the first place? We'll start by discussing why the above transitions and the relative short history of the field might actually provide a case against scientific realism and then suggest that realism should still be the default philosophical position for generative linguists based on both the success of the framework and the initial reasons for cognitive revolution in science. In what follows, I am not going to discuss the general philosophical reasons for and against scientific realism in the philosophy of science but rather my focus will be on those reasons which are relevant to generative linguistics (for a more general account see Rowbottom 2019).

Scientific realism is the position that the elements and posits of our scientific theories are literally or at least approximately true of the natural world (see Boyd 2010). Put in another way, realists hold that our best scientific theories say something true of both the observable and unobservable worlds. Thus, what geologists quantify over in their theories - which often shares its ontology with commonsense views on objects - is equally as *real* as highly theoretical entities such as quarks and electrons (even before they were observable). Another important aspect of realism is the idea that the objects posited by our theories are mind-independent or observer independent.[17] This of course serves to mitigate traditional metaphysical scepticism and idealism as it establishes independent belief in the external world. As van Fraasen puts it,

> Science aims to give us, in its theories, a literally true story of what the world is like; and acceptance of a scientific theory involves the belief that it is true. This is the correct statement of scientific realism. (1980: 8)

Despite his certainty, nonidentical (and sometimes incompatible) definitions of scientific realism abound. So much so that Chakravartty (2011) claims that "[i]t is perhaps only a slight exaggeration to say that scientific realism is characterized differently by every author who discusses it". Nevertheless, one of the main pulls of the position is its explanatory power or rather an appeal to the best explanation. Specifically, the idea is that scientific realism is the most "natural" explanation of the predictive success of the enterprise as a whole. There are various versions of this basic idea, the most famous of which is the "no miracles argument" (Putnam 1975) or the claim that without scientific realism, the empirical success of science would be miraculous.[18]

[17]Hence the furious debates around the foundations of quantum physics and Heisenberg's uncertainty principle.

[18]Of course, van Fraasen (1980) and others (such as Wray 2007) disagree with very idea that this empirical success is in need of explanation or at least not the kinds of explanations realists provide. Laudan (1981)

Naturally, during its relatively short history, generative linguistics has achieved its fair share of empirical success. From the discovery of cross-linguistic patterns and principles to the explanation of movement, anaphora, island effects and countless other findings to a plausible account of the evolution of the faculty of language itself. A scientific realist explanation of these accomplishments would mean that linguists could assume a level of truth to their theories. In fact, many arguments for taking a realist attitude towards a theory involve the putative success of that theory (see Ladyman 2011).

There is, however, a serious worry lurking in the present setting. The short history of the field and its radical shifts might actually go in the opposition direction and militate against a scientific realist stance in generative linguistics. The idea is something of this sort: in many cases scientific realists want their claims to be understood in terms of mature sciences such as physics and chemistry where methodological practices have stabilised somewhat. But the phenomena I appeal to in the present work, such as rapid theory change, could be taken as evidence against the idea that linguistics is a *mature* science and thus an appropriate context for a realist position. Indeed, many generativists have openly remarked on the incipient nature of the field. In discussing the conceptualist or mental realist framework, Higginbotham avows that "strong conceptualism is at the present state of scientific knowledge not so much an *indefensible* position as an *inarticulate* one" (1991: 559). Such ruminations might be characteristic of a pre-paradigmatic stage of scientific development in Kuhnian terms.[19]

My response to this worry is related to mentalism and the original role of linguistics in the cognitive revolution of the mid-twentieth century. Part of the goal of the establishment of cognitive science, which was a cross-disciplinary project involving linguistics, cybernetics, information theory, early cognitive neuroscience and other fields, was to counter the influence of Behaviourism on psychology and the study of mind (see Miller 2003). In order to do this, the liberation of 'the mental' as a legitimate object of scientific inquiry needed to take place. Generative linguistics at the time was a leader in this nascent undertaking. As mentioned previously, the formal mathematical tradition in generative linguistics can be argued to have been established by Chomsky (1956) and (1957) respectively. Mentalism in linguistics, on the other hand, has a slightly different trajectory which can be traced back to Chomsky's (1959) review of B.F. Skinner's *Verbal Behavior* and returned to in Chomsky (1965) with the idea of linguistics as a study of *ideal* mental competence in a language (and

himself argues that is is possible to have approximate truth without empirical success, as well as successful reference without empirical success.

[19] I thank an anonymous reviewer for pointing this worry out to me.

linguistics as a subfield of psychology).

But importantly, for linguistics to lead the charge against behaviourism, the concept of language as a mental object or state needed scientific validity. Neither instrumentalism nor constructive empiricism would have adequately done the job. As Pylyshyn (1991) says of the time:

> [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*. (1991: 232)

Mental realism, or simply mentalism as it is often called, amounts to scientific realism about the object of linguistic theory. Generative linguistics aims to describe the true nature of language as mental competence and acceptance of that theory involves believing in its truth (to paraphrase van Fraasen above). Thus, although the theory might have undergone changes over time, the Kuhnian paradigm was established in part by its adherence to and immense success with relation to the larger cognitive scientific project which itself essentially takes mental states and objects to be real features of the world.[20] As Chomsky (1983: 156) himself states:

> "[A] mentally represented grammar and UG are real objects, part of the physical world, where we understand mental states and representations to be physically encoded in some manner. Statements about particular grammars or about UG are true or false statements about steady states attained or the initial state (assumed fixed for the species), each of which is a definite real-world object, situated in spacetime and entering into causal relations."

Maintaining a realist stance has, therefore, been of paramount importance to the movement in general.[21] In addition to arguments from the success of the field, without mental

[20]Another reason one might favour a paradigmatic understanding of generative linguistics is provided in Tomalin (2010) who adapts Kuipers' (2007) taxonomy of scientific research categories. At the top are (1) research traditions, e.g. generative linguistics itself (including phonology, syntax etc.), which are instantiated by (2) research programmes such as generative grammar (further subdivided into Standard and Extended Standard Theory, Minimalism etc.) or the parallel architecture which in turn have (3) core theories (such as the autonomy of syntax or recursion) and finally (4) specific theories of particular phenomena which share core theoretical tenets. "This seems reasonable since the phrase 'generative grammar' is standardly used to refer to different theories of generative syntax that have been proposed during the period 1950s-present, and, given this, it would be misleading to classify GG as being simply a 'theory' (Tomalin 2010, 317)".

[21]There is an interesting possible connection here between what Shapiro (1997) calls "working realism" in which mathematicians act *as if* some sort of Platonism is true (or even *should* do so) and the case of generative linguists who assume that some sort of mental realism is true. Of course such a position would be too weak to defeat anti-realism. I thank an anonymous reviewer for suggesting the connection.

realism, the status of generative linguistics as a cognitive science is uncertain and thus scientific realism is one of the core tenets of the paradigm in general.[22] Furthermore, theoretical linguistics has had ties with linguistic pathology or aphasiology for decades. This link provides another reason to think that aspects of linguistic theories are instantiated in mental structures. Consider so-called Broca's agrammatic aphasia which usually occurs following a lesion in the Broca's area in the left hemisphere of the brain. The deficit causes individuals to lose their ability to produce syntactically well-formed sentences but in many cases semantic and phonological ability remain intact. The autonomy of syntax has been a long-argued for position in generative linguistics and there seems to be some evidence of its cognitive reality here. Similarly, various disorders offer indirect confirmation of other theoretical linguistic distinctions such as those between open and closed class categories, e.g. telegraphic speech in which function words are omitted (see Gabig 2013).

But standard scientific realism won't do. Accepting that would endanger the successes of the paradigm *via* PMI-related difficulties mentioned above. So given that commitment to realism is a significant aspect of understanding generative linguistics, rendering antrealism inimical, and given that realism *simpliciter* is problematic in the ways suggested above, I propose that structural realism be adopted as not only a means of obviating the PMI but also accessing the true nature of a structural enterprise such as generative linguistics.

5 Structural Realism in Linguistics

The previous sections showed a theory in flux with each new stage seemingly jettisoning the achievements of the last. In such a scenario, the PMI seems especially problematic. Not only this but as mentioned before, the situation in linguistics is unique since practitioners of each epoch of the theory can still be found working within the remit of their chosen formalism. In section 2, I described some of the theoretical shifts in the generative paradigm since the 1950s. In section 3, I described the underlying mathematical formalisms utilised in service of the changing theory at each junction. (While in section 4, I provided an argument in favour of scientific realism about generative grammar). In this

[22] What I find interesting is that despite questioning the cognitive scientific link, many of the philosophical critics of generative grammar have similarly insisted on realist interpretations of their views. Katz and Postal (1991) move from talk of Platonism to describing their view as 'Linguistic Realism'. Devitt (2006) too considers himself a non-mentalist realist (in a more nominalistic sense). Scientific realism seems to be a commonly held position within the foundations of linguistics across the philosophical spectrum.

section, I want to use a structural realist analysis of linguistics to show that despite the former, the structures of the latter remained relatively constant or at least commensurable.

What is structural realism? One way of thinking of it is as the 'best of both worlds' strategy for dealing with PMI. Realists, as we have seen, have trouble holding on to the objects of their theories once better theories come along. Anti-realists, on the other hand, have trouble accounting for the unparalleled predictive and explanatory success of theories (whose objects don't refer to objects in reality). Structural realism offers a conciliatory intermediary position between these choices. Ladyman (1998: 410) describes the position as follows.

> 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).

There are two versions of structural realism in the philosophy of science. The first, initially proposed by Worall (1989), is epistemic in nature. The second, championed by French (2003), is an ontological proposal. The former involves the idea that all we can *know* is structure, while the latter is a claim about all *there is*. In other words, what is preserved across theory change is a kind of structure posited by the underlying equations, laws, models or other mathematical representations of the theories. Part of the reason I opt here for ontic structural realism is that there is an ontological component to PMI as mentioned before. Thus, we are not only interested in what is communicated or epistemically accessible between different theories over time but what these theories say exists as well. Both versions agree on the existence of structures. Where they differ is on their respective treatments of objects. Ontic structural realism takes an anti-realist stance here while the epistemic variety is agnostic. Thus, ontological answer to PMI is therefore that if we cannot be realists about the objects of our scientific theories, we can be realists about the structures that they posit.[23]

[23] At this point, one can glean how such a picture might enter into the debate concerning the ontological foundations of linguistics mentioned earlier. Unlike Platonists who claim among other things that languages are individual abstract objects like sets or mentalists who claim they are psychological or internal states of the brain, a structuralist might argue that languages are complex structures in part identified by abstract rules and physical properties. See Nefdt (2018) for a similar view.

From here, it is not hard to see what the argument of the present section is going to be, namely that different generations of generative grammar display structural continuity notwithstanding variation in theoretical commitment. The means by which we can appreciate this continuity is by considering features of the mathematical representations employed during the course of history which could affect my proposed analysis. Moss (2012: 534) has a similar idea when he discusses the contribution made by mathematical models to linguistic theory.

> [L]anguage comes to us without any evident structure. It is up to theoreticians to propose whatever structures they think are useful [...] Mathematical models are the primary way that scientists in any field get at structure.

In the previous section, I told a story about how the proof-theoretic grammars of the Standard Theory were transformed into X-bar representations which eventually led to the Merge operation in Minimalism. However, a remarkable fact about the structural descriptions generated by these various formalisms is that they share a number of essential features, (1) they generate the same sets of sentences (also called 'weak generative capacity'),[24] (2) they take a finite input and generate an infinite output, and (3) they can be represented hierarchically through tree structures (not to mention actual structural similarities such as the way in which Merge joins two independent clauses and the way it was proposed in early transformational grammar). None of these latter properties are trivial. For instance, dependency grammars can be shown to be weakly equivalent to phrase structure grammars but are represented by means of flat structures. Model-theoretic grammars, such as Head-Driven Phrase Structure Grammar, are usually hierarchically represented and can generate the same sets of sentences but do not have any cardinality commitments. In other words, these features are *preserved under various transformations* of linguistic theory (a particular means of identifying structural identity, see next section).

Before I move on to a discussion of what structural properties could be and how to identify structures within linguistic theory, it is important to note that there were a number of formal shifts present in the transitions from transformational grammars to Merge. I

[24] In fact, these equivalences go beyond the generative grammars. Minimalist Syntax (or the Stabler 1997 version), Phrase-Structure grammars, Tree-substitution grammars, Head-Driven Phrase Structure grammars, and Dependency grammars have been shown to share weak generative capacity. See Mönnich 2007. Contrast this with 'strong generative capacity' in which a grammar assigns the same structural descriptions, e.g. Context-Free Grammars (CFGs) assign trees to each sentence. Thus, dependency grammars are not strongly equivalent in this sense to CFGs (or phrase-structure grammars) since they assign rooted acyclic graphs to sentences and not rooted binary trees.

have already mentioned the top-down to bottom-up change and argued that from a structural point of view, this is largely a metaphorical distinction. There is, however, another property of formal representations of syntax which also shifted from early to later generative grammar, namely from derivational approaches to representational or constraint-based ones. Simply put, derivational approaches follow the proof-theoretic model discussed earlier, where given a certain finite input and a certain set of rules, a particular structured output is generated. Constraint-based formalisms operate differently. Rather than 'deriving' an expression as output from a rule-bound grammar, these formalisms define certain conditions upon expressionhood or what counts as a grammatical sentence of the language.

Chomsky discusses this shift in thought in the following way.

> If the question is real, and subject to inquiry, then the [strong minimalist thesis] might turn out to be an even more radical break from the tradition than [the principles-and-parameters model] seemed to be. Not only does it abandon traditional conceptions of ''rule of grammar" and ''grammatical construction" that were carried over in some form into generative grammar, but it may also set the stage for asking novel questions that have no real counterpart in the earlier study of language (Chomsky, 2000: 92).

Indeed, with the Minimalist agenda and the Merge operation, more constraint-based grammar formalisms were embraced and adopted. This latter approach contains a different idea of 'rule of grammar' and indeed 'grammar construction'. The formal difference can be understood in terms of how each type of formalism answers the so-called 'membership problem'. Decidability is an important aspect of formal language theory. Given a string w and a formal language $\mathscr{L}(G)$, there is a finite procedure for deciding whether $w \in \mathscr{L}(G)$, i.e. a Turing machine which outputs "yes" or "no" in finite time. In other words, a language $\mathscr{L}(G)$ is decidable if G is a decidable grammar. This is called the membership problem. What determines membership in a traditional proof-theoretic grammar is whether or not that string can be generated from the start symbol S and the production rules R. In other words, whether that string is recursively enumerable in that language (set of strings).[25] What determines membership in a constraint-based grammar is

[25]As pointed out to me by an anonymous reviewer, semi-decidability would work for recursive enumerability as well. For instance, first order logic is not decidable but its validity is recursively enumerable (although the complement of the validity problem, i.e. determining whether a given formula ϕ is not valid, is not recursively enumerable).

whether the expression fulfils the constraints set by the grammar (which are like axioms of the system). 'An 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). As mentioned above, GPSG and HPSG are formalisms of the latter variety. While phrase structure grammars and tree adjoining grammars fall within the former camp.

The interesting fact for our purposes is that Merge and Minimalism represent the fruition of the gradual shift from derivational grammars to constraint-based ones. However, Chomsky (2000) does not initially put much stock in this formal transition despite the strong statement quoted above. He considers the old derivational or 'step-by-step procedure for constructing Exps' approach and the 'direct definition... where E is an expression of L iff...E..., where ...-... is some condition on E' approach to be 'mostly intertranslatable' (Chomsky, 2000: 99).[26] Here he holds these formalism-types to have few empirical differences, I will consider this thought in more detail in the next section.

From a mathematical point of view, the same formal languages and the structures of which they are composed are definable through both generative enumerative and model-theoretic means. Traditionally, the formal languages of the Chomsky Hierarchy were defined in terms of the kinds of grammars specified at the beginning of the previous section. However, there are other ways of demarcating the formal languages without recourse to generative grammars. For instance, they can be defined according to monadic second order logic in the model-theoretic way. Büchi (1960) 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) then 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).

The point is that the same structures can be characterised by means of proof-theoretic or model-theoretic techniques. Thus, the move from the former to the latter should not be seen as a hazard to the structural realist account of linguistic theory I am proffering here. In fact, in the next section I hope to show that this situation provides strong support for this particular analysis of the history and philosophy of linguistics.[27]

[26]He goes on to 'suspect' that the adoption of the derivational approach is more than expository and might indeed be 'correct'.

[27]This scenario is guaranteed by Beth's theorem which states (of classical logic) that a non-logical term T is *implicitly* defined by the theory (or generated by the rules) iff an *explicit* definition of the term is deducible from the theory (as in the case of constraint-based or model-theoretic grammars). This effectively connects

Lastly, the analysis suggested here dovetails naturally with other proposals to extend the purview of structural realism beyond physics and chemistry. For instance, Kincaid (2008) discusses the possibility of such an analysis of the social sciences. He argues that for structural realism to be successful vis-a-vis the social sciences, it needs to be shown that 'social scientists talk about structures and not individuals' (Kincaid, 2008: 722) and that when such talk occurs 'the individuals do not matter and the structure does' (724). In other words, social theories which emphasise 'roles' and 'relations' over and above the individuals occupying those roles or standing in those relations count in favour of a structural realist analysis. Kincaid offers three cases which meet the aforementioned condition, (1) general claims about social structure (e.g. organisations, classes, groups etc.), (2) the cases of causal modelling (and a reinterpretation of the problem of 'underidentification'), and (3) equilibrium explanations (involving the relations between self-consistent variables).

Similarly to these cases in the social sciences, linguistics (especially syntax) provides examples of structure trumping individuals. There are a number of examples in syntax, the most stark of which is the positing of covert material or items based, I argue, purely on structural considerations. Covert material in syntax refers to elements of the derivation that receive no phonological spell out. In other words, they are unpronounced items licensed only by the fact that the syntactic analysis requires a certain role to be played. Simple cases involve the EPP principle mentioned above (where a language can posit a 'null subject' to fulfil the structural requirement) and DPs or determiner phrases which need not contain actual determiners (such as *a(n), the, every* etc.). Another example is the PRO postulate in syntax. This element is an entirely *null* noun phrase (or empty category) which means it too goes unpronounced phonologically. This analysis figures in infinitival constructions in which PRO is said to operate as the subject of infinitives, *Mary wanted John [PRO] to help her*. The behaviour of this structural element PRO is different from that of general anaphors, referring expressions, and pronouns, which means it gets its own category despite not being visible to surface syntax. The idea is that something needs to fill the role in order for the overall structure to work, and thus PRO is postulated.

For a more developed example consider the generative analysis of negation below. 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

---

the proof theory of the logic to the model theory. I thank an anonymous reviewer for directing me towards the applicability of Beth's theorem here.

language) and the other from Spanish (a negative concord language).

(1) I didn't not go to work today.

DN: I went to work today.[28]

(2) María no puede encontrar a nadie

Maria *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* [my emphasis], 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 look negative but actually bear the [uNeg] feature. In other words, some elements which look negative on the surface can be semantically non-negative in reality.[29] Finally, it is argued that in grammatically justified situations, a covert [iNeg] can be assumed to c-command (or take scope over) 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 *Op* (which is often covert) and one or more overt non-semantically negative elements. Now consider an example from Czech in which it is argued that no overt negative elements are at play in the negation at all!

(3) Dnes nikdo *ne*volá

NC: Today nobody calls

[Dnes$Op_{\neg[iNEG]}$[$_{TP}$$nikdo_{[uNEG]}$ $nevola_{[uNEG]}$]]

[28]English 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.

[29]In addition, more technically, 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.

In (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).[30] The grammar then assumes a covert operator to generate the negative meaning. So the individuals words themselves do not generate the negative meaning but rather an unseen operator or item assumed purely for structural purposes fulfils this role.

Thus, in line with Kincaid (2008), linguistics can be shown to have cases (I would argue, many more than the social sciences) in which 'individuals do not matter' and structural considerations drive explanation. As he points out, there are general claims concerning structure, in our case phrases, X-Bar (as we've seen), trees, and operations on trees; specific cases of structural analyses such as negation and the general positing of covert structure; and even movement, an essential component of generative grammar across its time-slices, in which an item moves from one position in the tree to another, is not motivated by the individual nature of that item but the structural constraints on the grammaticality of the phrase or expression in which it is found.[31] Therefore, it would not be a stretch to consider linguistics, and syntax more so, to be a structural enterprise and thus amenable to a structural realist analysis.

Essentially, establishing that structural realism (whether epistemic or ontic) is a viable ontology for a series of theories requires two conditions to hold. The first is that they can be expressed structurally (in the sense of Kincaid). I have done so above for linguistics. The second is that their structures can be shown to be equivalent or isomorphic (or at least some related weaker structural relationship pertains). Section 3 made the case for the latter condition.

However, Kincaid's conditions might serve us well in motivating a general structural realist framework for a given science but it does not answer the question of what exactly is structurally preserved across specific theories. For this task, Ladyman's (2011) comparison between phlogiston theory and Lavoisier's oxygen theory is useful.

> Phlogiston theory subsumed the regularities in the phenomena above by categorizing them all as either phlogistication or dephlogistication reactions where

[30] The 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. But of course nothing here rests on the ultimate truth of this particular account, it is merely meant to show the overall structural thinking involved in generative analysis.

[31] The literature of WH-movement, for instance, is vast and can be found is almost all textbooks on syntax. Interestingly, for our purposes, the early trace theory is structurally identical to the later Minimalist copy theory of movement. The latter serves an additional theoretical purpose of limiting the proliferation of objects in the ontology such as the indices required for traces.

> these are inverse to each other. This is a prime example of a relation among the phenomena which is preserved in subsequent science even though the ontology of the theory is not; namely the inverse chemical reactions of reduction and oxygenation [...] The empirical success of the theory was retained in subsequent chemistry since the latter agrees that combustion, calcification and respiration are all the same kind of reaction, and that this kind of reaction has an inverse reaction, and there is a cycle between plants and animals such that animals change the properties of the air in one way and plants in the opposite way. (99)

Here he suggests that phlogiston theory meets a commitment of structural realism (both epistemic and ontic) in being a case of the "progressive and cumulative" nature of science and "the growth in our structural knowledge of the world goes beyond knowledge of empirical regularities" (Ladyman, 2011: 98). Similarly the trace theory of movement although replaced with the copy-theory retains this structural knowledge of how to account for movement (cf island effects in section 3). If we follow the analogy with phlogiston, neither phlogiston nor traces have a reference to anything in the world but the structural strategies employed by the earlier theories were empirically successful to a certain extent and thus retained in the later ones.

The above case is a relatively clear example. Other cases are not as transparent. Consider again the move from phrase-structure grammars to the X-bar schema to *merge*. It is not obviously the case that the same structures are preserved across formalisms, at least not without additional stipulations. Phrase structure grammars, for instance, do not inherit their categories or function from their parts as is the case with X-bar theory. This property is called endocentricity (as we saw in section 3). In X-bar theory, a sentence (previously S - exocentric) is taken to be an Inflectional Phrase projected from the verb (endocentric). You can capture this property with merge but only by means of labels. Headed constructions (endocentric) can be and are represented in many phrase structure rules. However, they are not essentially endocentric. Rather linguists have traditionally restricted themselves to the endocentric formulations implicitly. Whereas the X-bar formalism is makes this property explicit. Consider the rules for NPs, VPs, PPs below:

i $NP \rightarrow Det, N$

ii $VP \rightarrow V, NP$

iii $PP \rightarrow P, NP$

In fact, NPs are considered DPs or determiner phrases now as the head is thought to be the determiner. Besides the explicit endocentricity of X-bar theory, the formalism also showed that specific rules can be generalised to structures involving the categories of SPEC, Head, and Comp, across the board. In other words, all of the rules from (i) to (iii) (and many more) can be simply captured by either of the structure rules shown in section 3 during the discussion of X-bar.

Thus, the "progressive and cumulative" growth in our structural knowledge is based in the realisation of the generalisability of headed constructions and projection. A structural feature inherited by *merge* (with labels) in an even more abstract manner (as shown in section 3).[32]

Nevertheless, without a more precise notion of structure or structural property, the analysis only serves to illuminate structural similarity. The last section aims to make more precise the notion of structure at play and in general how structural comparisons can be achieved.

6 Structural Properties and Linguistic Analysis

The last aspect of this account of the scientific nature and history of linguistics will involve a brief detour into the ontology of structures themselves. In so doing, I hope to suggest a particular path, in line with a proposal from Johnson (2015), for how linguists might identify the relevant structures of their science, especially with relation the PMI.

What is a structure? The most common definition found in the literature is the set-theoretic one. "A structure S consists of (i) a non-empty set U of individuals (or objects), which form the domain of the structure, and (ii) a non-empty indexed set R (i.e. an ordered list) of relations on U, where R can also contain one-place relations (Frigg and Votsis, 2011: 228). Another term for such structures is "abstract structures" which means that both the objects in their domain of U and the relations on R have no material content (i.e. they need not be interpreted). Although the set-theoretic notion is commonplace, it remains controversial. Landry (2007) convincingly argues that different contexts require different structures (Kincaid (2008) similarly argues for a case by case application of structural realism). Muller (2010) rejects both the set-theoretic and category-theoretic (see Awodey 2006) account in favour of an entirely novel approach. And a number of others propose alternative frameworks such as the graph-theoretic approach of Leitgeb and Ladyman (2008).

[32] I thank an anonymous reviewer for pressing me on this point.

Since directly defining structures can be a fraught exercise and ultimately "[a] structuralist perspective is one that sees the investigation of the structural features of a domain of interest as the primary goal of enquiry" (Frigg and Votsis, 2011: 227), another path to grasping structures might be through the related notion of a structural feature or property. There are at least two possible ways in which to identify structural properties in the literature, one in terms of direct *definability* and another *via* a particular notion of *invariance* across structures. Intuitively, the first kind of characterisation relies on the internal relations of a given formalism. For instance, what identifies the structure of the natural numbers are the axioms of Peano arithmetic interpreted either through the Zermelo numerals or von Neumann ordinals (which have distinct properties). On the invariance account, there is some process of abstraction across similar systems of relations and the homing in on the invariant aspects. Thus, it might involve identifying whatever is true of or held constant across isomorphic systems or somewhat more formally 'structural properties of objects in a system *S* are specified here as those properties that the objects 'keep' when making isomorphic copies of *S*' (Korbmacher and Schiemer, 2018: 305).[33] In the case of linguistics, isomorphisms might be too strong, however. Homomorphisms or weaker structural mappings might also identify invariant structure for our purposes.

There are reasons in favour of both options. For instance, Nefdt (2018) opts for the definability approach for linguistic structures in accordance with a noneliminative structuralist account of mathematical objects (Shapiro (1997)) (i.e. the idea that singular objects are retained in the overall structural picture). There the task was to provide a possible response to another infamous puzzle posed by Benacerraf (1973) concerning the ontology of abstract objects. However, one problem with using the same strategy for addressing PMI type worries is that comparison across structures is difficult on the definability view. If what determines the identity of linguistic structures are the internal relations of the grammars, then characterising structural continuity across generations of grammar formalisms with distinct internal relations (i.e. grammar rules) is hard.[34] Consider the operations of *substitution* in Tree Substitution Grammar (TSG) and *adjunction* in Tree Ad-

[33]One may be tempted to consider these to be two identical or converging ways of carving up the same turkey. But according to Korbmacher and Schiemer (2018), once we move from the informal to the formal characterisations of these concepts, their differences become more apparent. See below.

[34]Hard but not impossible. In his dissertation, Meier (2015) compares Bloomfield on *substitutes* and Harris' *kernel sentences* toward an intertheoretical account of structural continuity. He defines a metatheoretical notion of theory reducibility for this purpose (i.e. Bloomfield is reducible to Harris and Harris to early Chomsky). Thus, he shows that the internally defined aspects of a theory are amenable to structural analysis and comparison. In this case, however, he is limited to epistemic structural realism which takes no stance on the properties of the objects of the structures in question.

joining Grammar (TAG).[35] TSGs and TAGs are weakly equivalent, their internal operations are similar and the structural output (i.e. binary trees) is identical. But TSGs cannot deal with phenomena like adjectival modification as TAGs do (part of the reason for the latter's development). In other words, internal rules might look similar in terms of their structural output but be distinct in terms of the internal structures themselves. In fact, *ante rem* or noneliminative structuralism in general faces problems with interstructural identity for precisely this reason. So much so that one advocate of the theory considers it undefinable (Resnik) and the other opts for primitively defining it (Shapiro).

Formally, the definability account 'subsumes the invariance account' due to the fact that isomorphic systems are semantically equivalent (or the 'isomorphism lemma' in model theory). The invariance account, however, does not subsume the definibility one since 'it is not *generally* the case that invariant properties are also definable in the particular language of the theory in question' (Korbmacher and Schiemer, 2018: 314). In our case, this means that the invariance account can assist in the revelation of structural continuity over and above the specific internal rules of particular grammars.

In fact, following notions of symmetry and invariance in physics, Johnson (2015) sets the precedent for the adoption of invariance considerations in linguistics, albeit for different purposes. He starts by modifying Chomsky on the notion of 'notational variants' or the idea that 'two theories (formal grammars, etc.) are notational variants iff they are empirically equivalent' (Johnson, 2015: 163) or following Chomsky (1972) do not differ in empirical consequences. He then presents a compelling case for applying a measure-theoretic analysis to generative linguistics. But before doing so he makes a few interesting points which verge on a structural realist view without endorsing (or mentioning) the possibility.

> Collectively, the notational variants of a theory determine the empirically 'real' or 'meaningful' structure of any one of the theories taken individually. This meaningful structure is often not identifiable without recourse to notational variants (i.e. symmetries) (Johnson, 2015: 164).

He goes on to claim that notational variants can shed light on which parts of theories are of empirical consequence and which parts are mere artifactual structure. For instance, consider the difference between two ways of representing temperature, Celsius

[35] *Substitution* involves replacing a non-terminal leaf in a tree with a new tree whose root node is labelled with the same non-terminal in order to create larger trees. *Adjunction* allows insertion of auxiliary trees within larger trees at various points. TAGs incorporate both mechanisms. See Rambow and Joshi (1997) for more details.

and Fahrenheit respectively. The 'real' empirical content or structure of temperature is determined by their convergence or intertranslatability. Anything *sui generis* about either system of representation is merely artifactual.

For a more controversial case involving linguistics, consider the discrete infinity postulate of generative grammars. If certain model-theoretic treatments of syntax do not entail cardinality properties (are 'cardinality neutral', see Pullum (2013)), then discrete infinity is an artifact of the formalisms used not a real feature of linguistic structure (see Nefdt (2019a) for a related argument). Johnson identifies the 'invariance principle' which roughly states that what is interesting empirically about a given formal grammar is not what it says but rather what it agrees with every other grammar on. This principle might be useful for providing an answer to the problem in Quine (1972) related to the psychological plausibility of multiple equivalent grammars, which is one target of Johnson's account, but in its strong form it also militates against a notion of scientific progress across generations of formal grammars. Thus, I would argue that certain so-called 'artifactual' or non-invariant structure can actually shed light on the differences and potential progress of later formalisms.

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 *vice versa*) and this can lead to predictions based on logical extensions of the mathematics or merely interpreting 'unused' mathematical structure.

Perhaps this is just to say that invariance is not the only means of identifying structural features relevant to understanding theory change (definability could also prove to be of ancillary use). Nevertheless, it is a useful concept for identifying those parts of linguistic theory that have remained constant and those parts that have changed in a commensurable manner.

Before concluding, it is important to address one residual issue related to ontology. In section 4, I argued that mental realism provides and has traditionally provided a reason for generative linguists to be realists in some sense. In the subsequent sections, I developed a kind of realism I believe supports the true mathematical nature of the field. But what, then, are these structures or structural properties?

The immediate answer is that they are cognitive structures. Indeed this is plausible (especially in light of the 'argument from aphasia' discussed in section 4). Johnson (2015) implies as much given that his proposal was meant to target Quine's argument concerning the problem of mental reality of weakly equivalent grammar formalisms. Thus, if notational variants or equivalent mental grammars homed in on invariant structure, then presumably that structure is cognitive in nature. In this way, adhering to the mental realism of section 4 is compatible with the structural realism advocated in subsequent chapters.

However, once again, one might worry that mental realism might be the wrong ontological interpretation of generative linguistics. Devitt (2006), for instance, proposes a thoroughgoing nominalist ontology which aims to interpret the field and its successes. Platonism is another infamous option, despite its few adherents. While Santana (2016), Nefdt (2018), and Stainton (2014) all proffer pluralist alternatives. On the latter's view, as an example, languages are hybrid ontological objects with part mental, part abstract, and part social structure. He states his position 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).

He offers two general arguments for his ontological pluralism. The first is similar in kind to Santana's (2016) proposal that various ontologies have important pieces of the puzzle of language to contribute and neglecting any would be tantamount to serious omission. The second crucially goes beyond this inclusivity to argue for compatibility. A detailed exposition of the overall view is beyond the present scope but it does offer a viable ontology for linguistics that does not obviously eschew the mental realist position standardly assumed. I believe that extending structural realism to this ontological picture would not be a particularly difficult exercise. The resulting structural realism would then pick out hybrid structural properties. Again, the details will have to be left for another occasion. For now, suffice to say, that although structural realism is compatible with mental realism or mentalism, it doesn't require that view and can be tailored to other metaphysical frameworks.

7 Conclusion

The primary goal of the paper was to argue that adopting the structural realist framework for linguistics has a number of philosophical advantages. Not only does it explain radical theory change in an anti-pessimistic manner but it also resolves the conflict over the ontology of natural languages in a way both consistent with accounts of the natural sciences and the formal motivations of the initial generative approach to the study of language. There are of course many further details necessary for a comprehensive defence of such an account, both historical and philosophical. This work serves to chart just one path toward the successful application of certain ideas in the philosophy of science to theoretical linguistics.

Acknowledgements

I would like to thank Geoff Pullum and Michael Kac for their insights into the historical and technical linguistic details as well as Steven French for offering guidance on the structural realist aspects of the article. I am also grateful to two anonymous reviewers of *Synthese* for their excellent comments on both the philosophical and technical elements of the work. All remaining errors are my own.

References

Awodey, S. 1996. 'Structure in Mathematics and Logic. A Categorical Perspective'. *Philosophia Mathematica* (3) 4. 209-237.

Biberauer, T., and Zeijlstra, H. 2012. 'Negative Concord in Afrikaans: filling a typological gap'. *Journal of Semantics*, 29: 345-371.

Bickerton, D. 2014. *More than nature needs: language, mind, and evolution.* Cambridge: MA: Harvard University Press.

Büchi, R. 1960. Weak second-order arithmetic and finite automata. *Zeitschrift für mathematische Logik und Grundlagen derMathematik* 6. 66-92.

Bueno, O., & Colyvan, M. 2011. 'An inferential conception of the application of mathematics'. *Nous* 45(2). 345–374.

Chakravartty, A. 2011. 'Scientific realism'. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. Stanford: Stanford University

Chomsky, N. 1956. 'Three models for the description of language'. *IRE Transactions on Information Theory* 2. 113–123.

Chomsky, N. 1957. *Syntactic structures.* The Hague: Mouton.

Chomsky, N. 1959. 'Review of Skinner's *Verbal Behavior*'. *Language* 35: 26-58.

Chomsky, N. 1964. *Current issues in linguistic theory.* The Hague: Mouton.

Chomsky, N. 1965. *Aspects of a Theory of Syntax.* MIT Press, Cambridge, Massachusetts.

Chomsky, N. 1970. 'Remarks on nominalization'. In R. Jacobs & P. Rosenbaum (eds.), *Readings in English transformational grammar*, 184–221. Waltham, MA: Ginn.

Chomsky, N. 1972. *Studies on semantics in generative grammar.* The Hague: Mouton.

Chomsky, N. 1975. *The logical structure of linguistic theory.* Netherlands: Springer.

Chomsky, N. 1983. 'Some conceptual shifts in the study of language'. In L. S. Cauman, I. Levi, C. D. Parson and R. Schwartz (eds.) *How many questions? Essays in Honor of Sidney Morgenbesser.* Indianapolis, Hackett.

Chomsky, N. 1986a. *Knowledge of language: its nature, origin, and use.* New York: Praeger.

Chomsky, N. 1986b. *Barriers.* MIT Press. Chomsky, N. 1995. *The Minimalist Program.* Cambridge, MA: MIT Press.

Chomsky, N. 1998. 'Minimalist Inquiries: The Framework'. *MIT Occasional Papers in Linguistics* 15.

Chomsky, N. 2000. 'Minimalist inquiries'. In R. Martin, D. Michaels & J. Uriagereka (eds.), *Step by step: essays on minimalist syntax in honor of Howard Lasnik*, 89–155. Cambridge, MA: MIT Press.

Chomsky, N. 2008. 'On phases'. In R. Freiden, C.P. Otero & M. Zubizarreta (eds.), *Foundational issues in linguistic theory*, 133–166. Cambridge, MA: MIT Press.

Cowie, F. 1999. 'What's within? Nativism Reconsidered'. Oxford: Oxford University Press.

Culicover, P. 2011. 'Core and periphery'. In P. Hogan (ed.), *The Cambridge Encyclopedia of the Language Sciences*, 227–230. Cambridge University Press.

Devitt, M. 2006. *Ignorance of language.* Oxford: Oxford University Press.

Freidin, R. 2012. 'A brief history of generative grammar'. In G. Russell & D. Fara (eds.), *The Routledge companion to philosophy of language*, 895–916. New York: Routledge.

French, Steven. 2006. 'Structure as a weapon of the realist'. *Proceedings of the Aristotelian Society* 106, 167–185.

French, S. 2011. 'Shifting the structures in physics and biology: a prophylactic promiscuous realism'. *Studies in History and Philosophy of Biological and Biomedical Sciences* 42. 164–173.

French, S. & Ladyman, J. 2003. 'Remodelling structural realism: quantum physics and the metaphysics of structure'. *Synthese* 136. 31–56.

Frigg, R., and Votsis, I. 2011. 'Everything You Always Wanted to Know About Structural Realism But Were Afraid to Ask'. *European Journal for Philosophy of Science* 1, 2011, 227-276.

Gabig, C. 2013. 'Telegraphic Speech'. In: Volkmar F. (ed.) *Encyclopedia of Autism Spectrum Disorders*, 125-139. Springer, New York, NY.

Horstein, N. 2009. *A Theory of Syntax: Minimal Operations and Universal Grammar*. Cambridge University Press.

Jackendoff, R. 1977. *X' Syntax*. Cambridge, MA: MIT Press.

Jackendoff, R. 2002. *Foundations of language: brain, meaning, grammar, evolution*. Oxford: Oxford University Press.

Jäger, G., & Rogers, J. 2012. 'Formal language theory: refining the Chomsky Hierarchy'. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367. 1956–1970.

Johnson, K. 2015. 'Notational variants and invariance in linguistics'. Mind & Language 30(2). 162–186.

Joseph, J. 1999. 'How structuralist was "American structuralism"? *Henry Sweet Society Bulletin* 33. 23–28.

Katz, J., & Postal, P. 2011. 'Realism vs. conceptualism in linguistics'. *Linguistics and Philosophy* 14(5), 515–554.

Kincaid, H. 2008. 'Structural Realism and the Social Sciences'. *Philosophy of Science*, vol. 75, No. 5, pp 720-731.

Korbmacher, J., & Schiemer, G. 2018. 'What are structural properties?' *Philosophia Mathematica*. Vol. 26, No. 3, Pp 295-323.

Kornai, A., & Pullum, G. 1990. 'The X-bar theory of phrase structure'. *Language* 66(1). 24–50.

Kuipers, T. 2007. *General philosophy of science: Focal issues*. Amsterdam/London: Elsevier/North Holland.

Ladyman, J. 1998. 'What is structural realism?' *Studies in the History and Philosophy of Science* 29(3). 403–424.

Ladyman, J. 2011. 'Structural realism versus standard scientific realism: the case of phlogiston and dephlogisticated air'. *Synthese* 180. 87–101.

Langendoen, T. 2003. 'Merge'. In A. Carnie, H. Hayley & M.Willie (eds.), *Formal approaches to function in grammar: in honor of Eloise Jelinek*, 307–318. Amsterdam: John Benjamins.

Lappin, S., Levine, R. & Johnson, D. 2000. 'The structure of unscientific revolutions'. *Natural language and linguistic theory* 18(3). 665–671.

Landry, E. 2007. 'Shared Structure Need not be Shared Set-Structure'. *Synthese*, vol. 158(1): 1-17.

Laudan, L. 1981. 'A confutation of convergent realism.' *Philosophy of Science* 48. 19–49.

Leitgeb, H., and Ladyman, J. 2008. 'Criteria of Identity and Structuralist Ontology'. *Philosophia Mathematica* 16, 388-396.

Lenneberg, E. 1967. *Biological foundations of language*. New York: John Wiley and Son.

Lobina, D. 2017. *Recursion: a computational investigation into the representation and processing of language*. Oxford: Oxford University Press.

Ludlow, P. 2011. *Philosophy of generative grammar*. Oxford: Oxford University Press.

Matthews, P. 2001. *A Short History of Structural Linguistics*. Cambridge University Press.

Meier, T. 2015. *Theory Change and Structural Realism. A General Discussion and an Application to Linguistics*. Dissertation, Ludwig-Maximilians-Universität München.

Mönnich, U. 2007. 'Minimalist syntax, multiple regular tree grammars, and direction preserving tree transducers'. In J. Rogers & S. Kepser (eds.), *Model theoretic syntax at 10. ESSLLI' 07 workshop proceedings*, 68–95. Berlin: Springer.

Moss, L. 2012. 'The role of mathematical methods'. In G. Russell & D. Fara (eds.), *The Routledge companion to philosophy of language*, 533–553. New York: Routledge.

Muller, F. 2010. 'The Characterisation of Structure: Definition versus Axiomatisation'. In: Stadler F. (eds) *The Present Situation in the Philosophy of Science. The Philosophy of Science in a European Perspective*, vol 1. Springer, Dordrecht.

Nefdt, R. 2018. 'Languages and other abstract structures'. In C. Christina & M. Neef (eds.), *Essays on linguistic realism*, 139–184. Amsterdam: John Benjamins.

Nefdt, R. 2019a. 'Infinity and the foundations of linguistics'. *Synthese*. Vol 196, No. 5, pp 1671-1711.

Nefdt, R. 2019b. 'Linguistics as a Science of Structure'. In McElvenny, J. (ed.) *Form and Formalism in Linguistics*, pp 175-196. Berlin: Language Sciences Press.

Newmeyer, F. 1996. *Generative linguistics: a historical perspective*. New York: Routledge.

Poole, G. 2002. *Syntactic theory*. Great Britain: Palgrave.

Postal, P. 1964. 'Constituent structure: a study of contemporary models of syntactic description'. *International Journal of American Linguistics* 30.

Pullum, G. 1983. 'How many possible human languages are there?' *Linguistic Inquiry* 14(3). 447–467.

Pullum, G. 2011. 'The mathematical foundations of Syntactic Structures'. *Journal of Logic, Language and Information* 20(3). 277–296.

Pullum, G. 2013. 'The central question in comparative syntactic metatheory'. *Mind & Language* 28(4). 492–521.

Pullum, G & Scholz, B. 2001. 'On the distinction between model-theoretic and generative enumerative syntactic frameworks'. In de Groote. P., G. Morril & C. Retoré (eds.), *Logical* *aspects of computational linguistics: 4th international conference*, 17–43. Berlin: Springer.

Pullum, G & Scholz, B. 2007. 'Tracking the origins of transformational generative grammar'. *Journal of Linguistics* 43(3). 701–723.

Putnam, H. 1975. *Mathematics, Matter and Method*. Cambridge: Cambridge University Press.

Pylyshyn, Z. 1991. 'Rules and representations: Chomsky and representational realism'. In A. Kasher (Ed.), *The Chomskyian turn* (pp. 231–51). Cambridge, MA: Blackwell.

Quine, W. 1972. 'Methodological reflections on current linguistic theory'. In D. Davidson & G. Harman (eds.), *Semantics of natural language*, 442–454. Dordrecht: D. Reidel.

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.

Rogers, J. 1998. '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)*.

Ross, J. 1967. *Constraints on variables in syntax*. PhD dissertation., MIT.

Shapiro, S. 1997. *Philosophy of mathematics: structure and ontology*. Oxford: Oxford University Press.

Rowbottom, D. 2019. 'Scientific realism: what it is, the contemporary debate, and new directions'. *Synthese*, 196: 451-484.

Stabler, E. 1997. 'Derivational minimalism'. In C. Restoré (ed.), *Logical aspects of computational linguistics*, 68–95. Berlin: Springer.

Stainton, R. 2014. 'Philosophy of linguistics'. *Oxford Handbooks Online*.

Stokhof, M., & van Lambalgen, M. 2011. 'Abstractions and idealisations: the construction of modern linguistics'. *Theoretical Linguistics* 37(1/2). 1–26.

Thatcher, J., & 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.

Tomalin, M. 2006. *Linguistics and the formal sciences*. Cambridge: Cambridge University Press.

Tomalin, M. 2010. 'Migrating propositions and the evolution of Generative Grammar'. In D. Kibbee (Ed.), *Chomskyan (r)evolutions* (pp. 315–337). Amsterdam: John Benjamins Publishing Company.

Van Fraasen, B. 1980. *The Scientific Image*. Oxford: Oxford University Press.

Worall, J. 1989. 'The best of both worlds?' *Dialectica* 43(1/2). 99–124.

Wray, B. 2007. 'A Selectionist Explanation of the Success and Failures of Science'. *Erkenntnis*, 67(1): 81-89.
Zeijlstra, H. 2004. *Sentential negation and Negative Concord.* Ph.D. dissertation, University of Amsterdam.