1 Introduction

1.1 What Is Theoretical Linguistics?

Defining theoretical linguistics, the remit of the present philosophical investigation, is a surprisingly fraught task. I say 'surprisingly' because one might expect that the domain is specifiable in terms similar to other fields prefixed with the same modifier. *Theoretical* physics is contrasted with experimental physics in terms of the kinds of methods and tools used in its exploration. For instance, theoretical physicists often make use of mathematical frameworks such as group theory to identify properties of symmetry or invariance in natural structures. They can incorporate sweeping idealisations in pursuit of laws of nature (Cartwright 1983). On the other hand, experimental physicists, such as those in big data cosmology, focus their efforts on applying statistical techniques to questions related to the origins of the universe. Experimental physicists more generally conduct real-world experiments (including simulations) to test and confirm theoretical posits or hypotheses. In a different vein, *theoretical* philosophy differs from *practical* philosophy in a shift in emphasis from abstract reality to practical, quotidian matters. In some cases, the distinction is captured by the difference between descriptive and normative contexts. Practical philosophy involves what we *ought* to do while theoretical philosophy aims to uncover what we *in fact* do (or *might* do in other possible worlds). Of course, this characterisation is overly simplistic.[1] There are fields that live in both worlds such as metaethics or mathematical physics. There are also fields where the distinction doesn't seem to hold such as biology and chemistry.

The nomenclature of philosophy and physics is equally unhelpful in the case of linguistics. Theoretical linguistics is indeed a descriptive enterprise, but so is experimental linguistics. In some cases, such as generative grammar, discrete mathematics characterises the methodological core of the practice. However, in others, such as probabilistic linguistics, continuous mathematics is favoured (Bod *et al*. 2003). Similarly, the status of experimentation is unclear in

[1] For one thing, the new conceptual engineering movement in philosophy aims at replacing or ameliorating 'defective' concepts. See Isaac *et al*. (2022) for an introduction to the new field.

https://doi.org/10.1017/9781009082853.002 Published online by Cambridge University Press

linguistics. Corpus studies are becoming more prominent in theoretical contexts with some theorists like Marantz (2007) even suggesting that the intuitions of linguists stand proxy for corpus data.

Methodology alone won't settle the target and scope of theoretical linguistics. Where normativity plays a role, it's unlikely to be one that distinguishes between theoretical and other fields of linguistics. Various theoretical approaches, from generative grammar to dynamic syntax, embrace different tools and methods, some formal and others empirical.

Haspelmath (2021) discusses a similar issue when he distinguishes between 'theoretical', 'general', and 'particular' linguistics. 'Theoretical', for him, cannot be contrasted with experimental since experimental work in linguistics often serves to push theory. He finds the distinction between 'theoretical' and 'applied' preferable since work in language pedagogy, automatic speech processing, and speech therapy set out to contribute to the resolution of practical problems and 'not necessarily in furthering theoretical understanding' (Haspelmath 2021, p. 4). But this is tricky: language pathology has a long history of informing theoretical pursuits. Asphasia studies or the general study of the linguistic effects of brain damage, for instance, have cemented theoretical distinctions and concepts like function versus content words, syntactic versus paratactic constructions, and the modularity of mental grammars.[2]

The distinction Haspelmath seems to be making is that the domain of theoretical linguistics (for both the general study of language itself and that of particular languages) is theory-driven in some fundamental way. But 'of or related to theory' assumes a neat dichotomy between theory and observation, which has been justly problematised in the philosophy of science. The more advanced the tools of observation, the more blurry the lines between observation and theory become. Think about the assignment of grammaticality to minimal pairs of sentences for a moment, either through introspection or corpus studies. Screening off syntactic well-formedness from semantic and pragmatic features is already a theory-laden activity. It presupposes autonomous syntax, which is a posit of a particular kind of theory (usually a generative one). Is an electron microscope a theoretical tool, enhanced observation, or both? Measurement is inherently theoretical. There's no clear distinction to be had between theory and observation, between science and facts.[3]

Despite the difficulty of the task, it's important to identify the field we're aiming to investigate. In some ways, the task has been made easier by the theoretical dominance of generative grammar in linguistics. Many excellent philosophical treatises have thus almost exclusively focused on it in their

[2] For a historical overview of the role aphasia studies have played in theoretical linguistics, see Elffers (2020).

[3] See Kukla (1996) for an argument that neither realists nor antirealists in the philosophy of science need to avail themselves of it.

https://doi.org/10.1017/9781009082853.002 Published online by Cambridge University Press

reflections: Newmeyer (1996), Ludlow (2011), and Rey (2020), to name a few. However, I don't plan to take generative grammar as metonymous for, or exhaustive of, theoretical linguistics. It's not the only game in town, nor was it ever. It's just one of many theoretical approaches to a distinct set of questions. Therein lies the clue as to my intended interpretation of the term 'theoretical linguistics'.

Theoretical linguistics is ultimately an explanatory project. The trouble is that the project has often been confused for its explanans, such as generative grammars or hierarchical tree structures. However, there are many and varied tools at the disposal of the theoretical linguist, well beyond these latter options. A better way to identify the project, in my view, is by appreciating its explananda, or the targets of its explanations. Thus, the way forward, as I see it, is to identify theoretical linguistics with a set of core theoretical questions.[4] These questions can be and are studied by means of numerous methods and approaches. They're unified not in approach but rather in their targets. The guiding set of questions is:

1a. *What is Language?*
 b. *What is a language?*
 2. *How do we acquire languages?*
 3. *How is linguistic communication possible?*
 4. *How did language evolve?*

To be a theoretical linguist, of whichever variety, you have to attempt to answer some, if not all, of these questions in a coherent manner. An applied linguist can get away without clear or scrutinised answers to the above sorts of questions. Of course, this interpretation doesn't constitute a necessary and sufficient set of conditions. I believe such a task would be largely fruitless. Nor does it cover every specific question a theoretical linguist might be interested in. We'll see more specific sub-questions in the following sections. It's my contention that they do all ultimately aim to produce answers to the general questions listed above. In Chapter 7, we'll come closest to views that diverge from this prescribed agenda. Computational approaches often have engineering goals in mind with human-level competence acting as little more than a benchmark, the so-called gold standard. Nevertheless, in keeping with the aims of this book, we'll still ask whether or not new approaches in artificial intelligence and computational linguistics can offer insights into the aforementioned theoretical questions. I'll argue that they do.

What's more important is that there's a distinct logical hierarchy in the list. The success conditions for any linguistic theory will depend on how one

[4] These are similar to Chomsky's (1965) nested adequacy conditions, with the addition of his later evolutionary bent and a focus on explaining communication, which is the focus of pragmatics and sociolinguistics.

answers (1), and the answer to (1) will determine the range of possibilities available for answers to (2) to (4). For instance, generative grammar (or biolinguistics) assumes that language consists of a modular mental system responsible for narrow syntax. Languages are specific settings of this system activated by varying external stimuli (from different language communities). Acquisition is largely explained by an innate module in the brain specific to our species. Communication is an exaptation. Language evolved for thought, and since it's mostly syntactic in nature, an operation like the set-theoretic one of 'Merge' can do the job of explaining its sudden emergence some 100 thousand years ago (Berwick & Chomsky 2016).[5] Of course, this particular sequence is negotiable at every turn. The package changes if we start with the idea that language is a mathematical representation of conventions in particular speech communities. David Lewis (1975) attempts to provide such a synthesis. Acquisition can essentially include sociolinguistic principles and discoveries with in its remit. Successful explanation along these lines would involve finding the correct model of the community's linguistic conventions.

However, methodologically, there can be convergence. Platonists, like Katz (1981), differ in their answer to (1) but insist that the methodology (of generative grammar) remains constant. In fact, they argue the methodology better fits their ontological paradigm (Postal 2003). In the next subsection, we move on to a discussion of one of the chief tools for answering these questions in the theoretical linguist's arsenal, namely, that of a *grammar*.

1.2 Grammar and Grammaticality

Most introductory linguistics textbooks start with a caution and a disclaimer: 'Prescriptivists keep out! Science ahead'. The idea is that the notion of a grammar is historically associated with various injunctions on writing and speaking 'properly' – 'Don't split your infinitives', 'Don't end sentences with prepositions', 'Avoid passives', and so on.[6] Students of a language need to disabuse themselves of these restrictive claims. The task of a grammarian is then not to prescribe arbitrary stylistic rules of 'proper usage' but to uncover the rules that govern *actual* usage. It's unclear whether or not linguists are solely interested in describing actual usage. Indeed, some corpus linguistics is directed at identifying patterns or regularities within various corpora, but for the most part, theoretical linguists see their tasks as inductive and ampliative. In so doing, they're invariably confronted with possibilia, or unactualised types of sentences, and constructions that are predicted and sometimes prohibited by the rules they

[5] We'll see more of this view in Chapter 2.
[6] See Pullum (2014) for a history of the 'fear and loathing' of the English passive, for instance.

https://doi.org/10.1017/9781009082853.002 Published online by Cambridge University Press

describe. What makes this process interesting is that this might be the point at which normativity creeps into the field (Kac 1994; Itkonen 2019; Pullum 2019).

In demarcating the space of acceptable or unacceptable strings of any language, one isn't only describing a state of affairs but prescribing certain legal and illegal operations. For example, if the only rule of English was that a Verb Phrase = Noun + Finite Verb, as in *Geoff sings*, then although hundreds of thousands of sentences would immediately be licensed, many other forms such as those involving determiners, adjectives, prepositions, and so on, would be banned or relegated to the inimical category of the 'ungrammatical'. Whether a grammar is a normative or descriptive device is a tricky philosophical question, one that often receives very little attention by linguists or even philosophers. Inferentialism in the philosophy of language and logic takes normativity to be central to the generation of meaning.[7] It enters linguistics via proof-theoretic semantics (Brandom 1994; Francez & Dyckhoff 2010; Peregrin 2015). Before we can approach this issue more broadly and what relation modern grammatical theory might even have to old-school grammar instruction (still alive and nitpicking in various popular writing books), we need to define what a grammar is and what role it plays in theoretical linguistics.

The 'orthodox' or mainstream generative view has it that 'grammar' plays multiple roles in the theory of language. Nevertheless, one overarching role, upon which Chomsky (1965, 1981, 2000) has repeatedly insisted, is that a grammar is a theory of a language, in the sense of a 'scientific' theory. For him, the target of theorising is our knowledge of language understood as a stable mental state of the language faculty. The overall job of linguistic theory is then to illuminate the structure of this knowledge or mental state. Specifically, there are two senses of 'grammar' common in the generative literature. The first kind of grammar attempts to map the contours of the mature state of the language faculty attained by an individual cogniser (her 'I-language'), while the second demarcates the settings of a deeper underlying universal patterning or the innate initial state of all language users. As Chomsky states, '[a]dapting traditional terms to a special usage, we call the theory of the state attained its *grammar* and the theory of the initial state *Univeral Grammar* (UG)' (1995b: p. 12). We'll see much more of these ideas in Chapter 2.

If linguistic theory is indeed scientific in any strong sense, then we might expect laws or regularities to emerge from our investigations. This expectation has led to at least two further trends in the field (that track the two senses above). The first and earliest has been the focus on syntax as a core aspect of the language faculty. One reason is that syntax is rather well behaved, math-

[7] There's a distinctive Wittgensteinian flavour to this framework, not only in the use-based theory but also in the later Wittgenstein views on mathematics as a 'network of norms' (Wittgenstein 1953, VII §67).

https://doi.org/10.1017/9781009082853.002 Published online by Cambridge University Press

ematically speaking. Early twentieth-century formal logic provided numerous insights into proof theory with the work of Post, Turing, Gödel, and Carnap, to name just a few prominent examples. Many of these results translate very well to the study of syntactic structures. In fact, formal language theory was invented as a subfield of linguistics that was directly informed by both logical structure and natural language constraints. Some early results, such as Chomsky (1956), aimed to show that natural language syntax outstripped the bounds of finite-state grammars and required context-free rules with transformations. Later work used data from languages such as Dutch and Swiss-German to show that context-free grammars were equally insufficient given the possibility of cross-serial dependencies (Shieber 1985). I'll return to some of these details in the next subsection, but for now, the basic idea is that syntactic complexity can be precisely characterised in terms of formal languages (generated by formal grammars). The trick is then to show that some natural construction, formally specified, exceeds the limits of a particular formal language by showing that it cannot be generated by the associated grammar. Here, particular patterns in particular languages inform the grammar qua scientific theory.

The second trend to emerge from the 'scientific expectation' was the search for linguistic universals, the ultimate regularities to be found in linguistic nature. Successfully identifying regular law-like patterns in cross-linguistic reality would go a long way to supporting the claim that grammars are theories of language that ultimately illuminate some sort of UG. If all languages, the world over, prescribe to a set of formally identifiable constraints, then studying these constraints might indeed reveal the underlying structure of language itself (Language with an uppercase 'L'). Despite decades of valiant attempts, languages (with a lowercase 'l') proved recalcitrant to such universal characterisation (see Evans & Levinson 2009 and Chapter 2).[8] The result was that more and more abstract properties were considered as candidates for universality. One popular such proposal is that *all natural languages are recursive* (Hauser *et al*. 2002). However, recursion is a formal property of grammar or representation. Iterative structures in natural language need not be represented as recursive (Lobina 2017). Furthermore, there's some question over where exactly recursive structure lies. As an aspect of the computational component of the UG, the claim becomes almost unfalsifiable since no particular language would offer counterevidence, whether or not it possessed the hallmarks of surface recursion like centre-embedding or propositional contexts like *Joan said that Irene believed that Angelika thought that ...*

[8] Some have followed Greenberg's 1963 infamous attempts at finding linguistic universals. Fascinating as this list is, it mostly comprises conditional patterns, many of which are not syntactic. Not to mention the initial sample was composed of around thirty of the world's eight thousand extant languages.

https://doi.org/10.1017/9781009082853.002 Published online by Cambridge University Press

This picture of the science of language can change with the role and definition of grammar. Some theorists such as Tiede & Stout (2010) and Nefdt (2016, 2019) have argued that grammars are more akin to formal models than scientific theories. The change of perspective has some profound consequences. For one thing, models are indirect representations of a target phenomenon. What this means is that grammars themselves (or some of their aspects) could be nonveridical. So instead of 'reading' the structures of the model as reflective of linguistic reality directly, one can appreciate a looser relationship between grammars and reality. For instance, the debate over the universal nature of recursion becomes a discussion over whether recursive structures are artefactual aspects of the model used or actual explananda of the system under investigation. This possibility has a knock-on effect on debates concerning the infinitude of natural language since it opens up the further possibility that talk of linguistic infinity is a mere simplification device similar to treating a complex system as *essentially infinite* in computer science even if it's in fact finite (see Savitch 1993; Nefdt 2019).

Another effect of this shift in interpretation involves the success conditions of grammars again. For instance, if formal semantics is in the business of assigning models to sets of sentences, then counterexamples would refute the models and require expansion (assuming nonmonotonicity). Theories can be more recalcitrant to contravening data, according to some linguists (see Chapter 7 for more).

To add to the complication, whatever our philosophical interpretation of grammar, the notion of 'grammaticality' can be somewhat detached from it. Interestingly, it seems less possible to detach grammaticality from the normativity debate, though. Assuming that grammaticality is a property of individual expressions or subexpressions of a language, what makes a sentence grammatical? The answers can converge and diverge whether you view grammars as scientific theories or models. However, determining the grammatical sentences does seem to involve deciding whether grammars are mental devices, reflections of community standards, or some hybrid of these and other options. For instance, if grammaticality is grounded in conventional practices of a linguistic community, then its application seems to be normative. Saying **I is hungry* is incorrect by the standards of the community since grammatical agreement between subject and copula is the norm.[9] Saying *Ek is honger* in Afrikaans is fine since that language has long since abandoned verb–subject agreement (and inflection from its parent language Dutch). If grammaticality is a property of formal expressions derived from a mental module or generative grammar, then certain violations might belong to the realm of performance

[9] Ungrammatical or unacceptable constructions are usually marked with an asterisk in the top-left corner of the sentence.

and not grammatical competence (see Chapter 7). You might say **I is hungry* because you're drunk, or trying to be funny, or both. We'll get to the competence–performance distinction later. But as Manning (2003) points out, this doesn't help when nonstandard usage is at play. Standard grammar rules tend to divide grammatical and ungrammatical strings absolutely (even if the grammar is undecidable).[10] The problem is that there are constructions and phrases that pop up all over human language (and corpora) that would be deemed simply ungrammatical in this strict generative sense (i.e. not generated by any discrete rule). Manning identifies one such construction, namely, *as least as*. This construction sounds strange at first but robustly appears across various texts. He claims that generative grammar (which he calls a 'categorical linguistic theory') is prescriptive in the sense that it places hard boundaries on grammaticality when these boundaries are much fuzzier in reality. Many theorists realise that grammaticality itself might be fuzzy. This possibility doesn't, however, rule out successfully using apparatus to tame it in discrete or binary terms. Historically, others have embraced the fuzziness and either advocated fluid grammatical catogories, or 'squishiness' (Ross 1973), or a fuzzy logic to capture it (Lakoff 1973). Others yet either reject grammaticality itself (Sampson & Babarczy 2013) or hope for theoretical illumination on the distinction between grammaticality and acceptability (Sprouse 2018), the latter more amenable to fuzzy characterisations than the former.

Of course, Chomskyans themselves can admit the possible gradable nature of grammaticality or acceptability at the performance level while arguing that discreteness is a useful idealisation nonetheless at the level of competence. Whether grammaticality is modelled discretely or continuously, deviations from the rules (or statistical generalisations) appear to elicit some normative force. The point generalises to deep issues about linguistic methodology going back generations. What do you do when your intuition-derived examples, and eventual laws, conflict with data from corpora? Deny the latter on pain of giving the descriptive game away?[11] Some American structuralists, like Charles Hockett, believed linguistics had to do both jobs at once: characterise corpora of utterances and explain unuttered possibilia. In a move inspired by Goodman and the normative theoretical device of reflective equilibrium (a bedrock in moral and political philosophy), Pullum claims the following of the epistemology of syntax:

> The goal is an optimal fit between a general linguistic theory (which is never complete), the proposed rules or constraints (which are not quite as conformant with the general

[10] Given a string w and a formal language $L(G)$, there's a finite procedure for deciding whether $w \in L(G)$, that is, a Turing machine that outputs 'yes' or 'no' in finite time. In other words, a language $L(G)$ is decidable if G is a decidable grammar. This is called the 'membership problem'. See Jäger & Rogers (2012).

[11] As we'll see, mainstream linguistics might co-opt the competence–performance distinction to avoid this issue altogether (see Chapter 2) or a 'Galilean' strategy in science (see Chapter 7).

https://doi.org/10.1017/9781009082853.002 Published online by Cambridge University Press

theory as we would like), the best grammaticality judgments obtainable (which are not guaranteed to be veridical), and facts from corpora (which may always contain errors). (2007, p. 37)

Theory development can then follow the Quine–Duhem thesis and the scientific holism it advocates. Confirmation depends on a tapestry of interconnected components, not individual linguists' judgements or speech corpora, wholly. Pullum offers this kind of picture as a response to Sampson's (2007) project aimed at both ridding linguistics of introspective data and divorcing grammatical theory from grammaticality entirely (along with the grammatical–ungrammatical distinction). In so doing, he unequivocally states: 'I take linguistics to have an inherently normative subject matter. The task of the syntactician is exact codification of a set of norms implicit in linguistic practice' (Pullum 2007, p. 39).

Grammar, grammaticality, normativity, and linguistic theory all seem to be interconnected. Before we return to those issues briefly below, we need to address one further (and related) tool that has proven powerful in the linguist's arsenal, that of formalisation.

1.3 Formal Approaches

Contemporary linguistics, as a discipline, is unique in many ways. One of the most interesting aspects of the field, one that sets it apart from many of the social sciences and humanities, is how highly formalised it is. This is apparent in syntax, which drew from work in proof theory, but phonology, semantics, and even pragmatics have all been modelled by formal apparatus of various kinds (e.g. optimality theory is one framework that has been used to formalise all of these subfields). Chomsky famously stated that:

> Precisely constructed models for linguistic structure can play an important role, both negative and positive, in the process of discovery itself. By pushing a precise but inadequate formulation to an unacceptable conclusion, we can often expose the exact source of this inadequacy and, consequently, gain a deep 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. (1957, p. 5)

Since then, precision has been the cornerstone of the enterprise.[12] But is there a reason beyond the rhetoric? Is formalisation more than just a tool in the language sciences? Let's evaluate both the 'negative and positive' sides of Chomsky's above claim.

12 Precision and formalisation were, of course, present in the linguistics done before Chomsky as well. Hilbert's programme in metamathematics greatly influenced early linguists like Bloomfield, Hockett, and Harris (see Tomalin 2006). But the formalism played a slightly different role later on, as we'll see.

There are, of course, many positive reasons for formalisation in the sciences more generally. Besides precision, formal theories tend to be explicit about the claims that are made. This feature in turn allows others to build on or critique those theories with more confidence. Mathematics has been 'unreasonably effective' at wrestling the hidden structure of the natural world into submission. This is especially true in physics where group theory, graph theory, and linear algebra have all proved successful in unearthing countless discoveries over the centuries (Wigner 1960). Philosophers, on the other hand, honed formal logic for the construction of their arguments and even used it to extract ontological consequences from their domains of inquiry (recall Quine's influential dictum: 'to be is to be the value of a bound variable' (1948)).

In linguistics, generative grammar was one of the disciplines at the helm of the classical cognitive revolution (Miller 2003). One of the core insights of this paradigm shift in the study of mind was the computational theory of mind, or CTM. The classical version of CTM proposed that the mind can be understood as a computational system or Turing machine of some sort. Chomskyan linguistics embraced not only the letter of CTM but also its punctuation. Formal grammars are formalised as recursive devices that enumerate potentially infinite sets via a finite set of rules. In fact, the $[\Sigma, F]$ or rewrite grammars of *Syntactic Structures* were modelled on post-production systems, Turing-complete recursive enumerators (see Pullum 2011). 'Each such grammar is defined by a finite set Σ of initial strings and a finite set F of "instruction formulas" of the form $X \rightarrow Y$ interpreted: "rewrite X as Y"' (Chomsky 1957, p. 22). A philosopher might recognise a similar procedure here to natural deduction in which you derive a certain formula or conjecture from an initial alphabet and the repeated application of the rules of inference. The field of formal language theory (FLT) became the dominant instantiation of Chomsky's statement at the start of this section. And although FLT has moved from mainstream linguistics to computational approaches (see Chapter 7), early results of the structure and complexity of natural language drew from it significantly.

Take, for instance, the proof of the context-freeness of natural language alluded to above. Without the formalisation of linguistic structure via formal grammars, actual proofs about the structure of language would have been impossible. A tempting thought might be that syntax can be well captured by means of a Markov chain or a simple statistical process involving initial states and transitions between them in sequence. This much is suggested by Saussurean structuralism.[13] Consider the finite state automaton (FSA) in Figure 1.1.

FSAs recognise regular languages (the least complex class of formal languages in the original Chomsky Hierarchy). In the diagram, $q0$ represents

[13] 'One of the principles defended by Saussure in the *Course in General Linguistics* is the principle of the "linear nature of the signifier" (1959, p. 70; 1916, p. 103), by which Saussure intends to say that words, like sentences, are concatenations of signs along a linear temporal axis (the time it takes to pronounce a word or sentence)' (Egré 2018, p. 670).

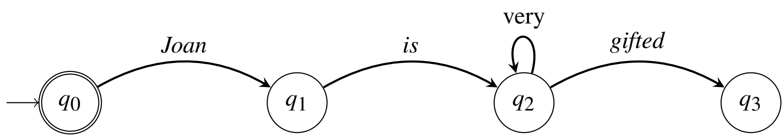

Figure 1.1 Finite state automaton

the input state while $q3$ represents the output (sometimes indicated by a full stop). This kind of formalism has produced significant results in information theory, and as a model of language, it would represent a behaviourist's dream: stimulus and response at its finest. Adding probabilities to the transitions serves as the basis of word-to-word sequencing and earlier predictive processing. A few interesting facts about FLT can already show us why human language cannot be represented by this kind of machine. Every automaton is an accepting machine. It can be paired with a particular formal grammar and language. Furthermore, what it accepts are sequences of strings of a certain complexity. It was argued in Chomsky (1956, 1957) that there are very common patterns in English that already outstrip the power of finite-state automata, namely, patterns of the $a^n b^n$ variety. Less abstractly, certain kinds of syntactic patterns show particular dependencies between units. *The man the woman saw left* fits the mould of the aforementioned pattern. The key to understanding why these sorts of dependencies, and the phenomenon of 'embedding' itself, cause so much difficulty is the appreciation that natural language syntax is a bit like a 'Matryoshka', or Russian nesting doll. Units aren't like Leibniz's *monads* (which were unique, indestructible substances making up the universe) but decomposable phrases that allow for further decomposition. The structural picture that goes with this idea is more like a tree than a line (see Figure 1.2).

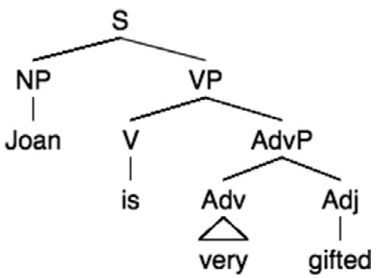

Figure 1.2 Tree for *Joan is very gifted*

The kind of grammar that produces these kinds of structures is phrase structure grammar.[14] The rule systems associated with them are 'context-free', which is a ring above the regular languages in terms of complexity. This means they can represent patterns that finite-state or regular grammars cannot,

[14] This class of automata is the pushdown automata with this neat feature called a 'stack' that acts like a little internal memory for the machine.

https://doi.org/10.1017/9781009082853.002 Published online by Cambridge University Press

especially the ones that have embedded constituents.[15] This makes them more expressive since they capture the patterns of the regular languages and more. So why don't we go for the most expressive grammars? In this case, these formalisms would give us the recursively enumerable languages accepted by Turing machines, basically, the machines that can compute any algorithmic computable function whatsoever. Well, it comes down to what the purpose of grammars were supposed to be, that is, they were supposed to generate *all and only all* the sentences of a language. Finite-state grammars don't get us the first 'all' and more complex grammars take us beyond the 'only all'. Now, of course, there's evidence to suggest that context-free or phrase structure grammars don't get us the first 'all' either (Shieber 1985). This argument, based on the syntax of Swiss-German, is a little out of the neighbourhood for us right now, and, in fact, Chomsky (1956) already anticipated that phrase structure grammars wouldn't do by themselves. They needed to be supplemented with transformation rules that could convert one kind of structure into another, for example, an active sentence into a passive.[16] Specifically, in transformational grammar, which characterised the 'Standard Theory' in generative grammar (Chomsky 1965), to get the syntactic structure of a sentence, you start with a 'kernel sentence' derivable by a context-free grammar or rewrite grammar (such as VP $\rightarrow$ V, AdvP above), then you apply syntactic transformations that insert, move, delete, or alter elements in the structures to which they apply.

This is all to say that formal grammars and formalisation have taught us, and can teach us, a lot about language at a very abstract level. But philosophically, it's even deeper than this. Remember, linguistics was a big player in the classical cognitive revolution (Miller 2003; Bever 2021). One of the core principles of this programme was that the mind is basically like a computer (at the software level). Thus, FLT, automata theory, and formal syntax all contributed to an understanding of what kinds of structures this mental computer was running on, specifically, when it processed language. This is part of the reason that linguistics started to look a lot like computer science. The other part of the reason is an insight attributed to the linguist Wilhelm von Humboldt (1836) that language 'makes infinite use of finite resources'. Chomsky, and others, have claimed that recursion theory (or computability theory) is the only way to make sense of this uniquely linguistic property. Hence, generative grammars with recursive rules.

Lastly, the computational analogy goes both ways. As Müller (2018, p. 6) notes:

15 Constituents are, again, those units of language that seem to act together. For example, they can usually be moved or fronted in English as in *that book, Irene coveted*. There are a number of other 'tests' for constituency common across syntax textbooks.

16 I've used phrase structure and context-free synonymously here. But there's reason to separate the notions; see Manaster-Ramer & Kac (1990).

> [a] further advantage of precisely formulated theories is that they can be written down in such a way that computer programs can process them. When a theoretical analysis is implemented as a computationally processable grammar fragment, any inconsistency will become immediately evident.

This last point tracks well with results in FLT that very often follow the pattern of showing why a particular grammar, formalism, theory, or structure isn't the right one, or at the right level, for particular purposes. Formalisation is a vastly important tool for theory building, but it can be even more useful as a debugger.

It's not only syntax that has undergone massive formalisation: semantics also benefited from the introduction of highly formalised models. If syntax drew inspiration from proof theory, then semantics looked to another side of logic, namely, model theory. Model theory is the mathematical study of interpreting structures or models of formal theories. Led by figures in philosophy like Richard Montague and David Lewis, formal semantics incorporated tools and insights from modal logic, type theory, and lambda calculus to model meaning as the functional complement to syntactic theory. Barbara Partee, the mother of formal semantics, is credited with saying 'lambdas really changed my life' (Partee 1996, p. 24).[17] This kind of mathematical life-changing experience is unsurprising in a field founded by a mathematician who infamously wrote:

> There is in my opinion no important theoretical difference between natural languages and the artificial languages of logicians; indeed I consider it possible to comprehend the syntax and semantics of both kinds of languages with a single natural and mathematically precise theory. (Montague 1970, p. 373)[18]

The road from formalisation to mathematisation to Platonism is paved with good scientific intentions. But despite the claims of a few philosophers and linguists,[19] most researchers in the language sciences embrace a different side of the field, one that aligns it less with logic and mathematics and more with fields such as physics, biology, and anthropology. It's to this face of linguistics that we now briefly turn.

1.4 Linguistics as an Empirical Science

One way in which to approach the question of the balance between formal versus empirical components in linguistics is by asking yourself this question: if all the mathematical facts about formal languages were settled, what would be

[17] Besides helping with naming complex set-theoretic objects associated with syntactic objects, lambdas aid in transforming syntactic objects like VPs into truth-conditional functions. We'll return to the tool of functional application in Chapter 6.

[18] Thomason (1974, p. 2) further claims that Montague held the view that 'syntax of English, for example, is just as much a part of mathematics as number theory or geometry'.

[19] Notably, Jerrold Katz (1981, 1996) and Paul Postal (2003, 2009).

left of linguistics? To hard-line platonists, the answer might be nothing. To most others, though, it's a panoply of possibility. Linguists study communication, phonetics, sign languages, and conventional regularities inherited from linguistic communities over time. Although some of the foundational tools have been borrowed from discrete mathematics, statistical and social scientific approaches have also played a role in some of the most important discoveries of the field.

Consider the famous case of class-based linguistic variation across societies. Intra- and inter-class regularities have been documented in speech patterns, register, and lexicon. For instance, '[Labov] soon established that when the collective pronunciation of a certain vowel changes without being noticed by the speech community, it is likely that upper-working class or lower-middle class women are leading the change' (Woschitz 2020, p. 163). His work on African American Vernacular English (AAVE) was groundbreaking. Chomskyans generally reject dialect studies as peripheral to linguistic theory and a nonstarter to scientific research on language. So-called external or E-languages are considered political more than scientific objects. But Labov showed that studying neglected (and oppressed) dialects of English, such as those spoken by working-class African American communities in the USA, provides a theoretical window into phenomena as diverse as structural variation, language contact, and language change. For example, Labov (1969) uncovered an interesting regularity that when copula contraction is permitted in Standard English, AAVE allows deletion (and the other way around). Furthermore, when contraction isn't possible, then neither is deletion. Consider the following:

1. Barbara said *she's* going/*Smart, that's what she's.
2. Barbara said *she* going/*Smart, that's what she.

His methodology wasn't mathematical. In fact, he used methodological innovations such as peer-group recordings to identify phonological as well as grammatical features across dialects and class structures. On a more systemic level, in a slightly more recent article, Labov (2010) argues that contrary to the contemporary trajectory of many oppressed languages toward extinction, residential segregation has reversed this fate for AAVE. Variability is key to understanding sociolinguistics. It doesn't deal in laws or mathematical certainty but contextually triggered phenomena, like when a linguistic environment triggers code-switching behaviour. Shifting from a job interview to a party with friends is likely to affect your register in systematic and measurable ways.

Contemporary biolinguists consider language to be a 'biological object' and claim to embrace the biological sciences as an alternative to formal logic and mathematics (we'll evaluate these claims more carefully in Chapter 8). Similarly to Labovian methodology, biology, for researchers like Berwick & Chomsky (2016), is more of a case study than the law-foraging expeditions of physics or mathematics. The regularities one discovers are often unstable and subject to change. Language is essentially a brain state, and our knowledge of neural circuitry is relatively nascent. Friederici *et al*. (2017) identify the Broca area (BA 44) as the brain centre of syntactic processing (as well as the dorsal pathway). '[W]ith respect to Broca's area, the activation of BA 44 as a function of syntax has been confirmed in many studies across different languages' (Friederici *et al*. 2017, p. 714). More specifically, they link these areas to the 'Merge' postulate of later generative grammar. Merge is an operation that takes syntactic objects and composes a labelled (unordered) set containing these objects, iteratively. Merge is meant to capture the allegedly universal property of the hierarchical structure of syntax, that is, sentences are composed of embedded phrases that can themselves be embedded, represented by a tree-like structure. The tools of biolinguistics are supposed to include fMRI studies, computed tomology, and other techniques used in event-related potential (ERP) research, in which direct brain responses are measured with specific stimuli. However, neuroscientists emphasise the property of neural plasticity, which militates against strong causal and structural claims about the brain (Silberstein 2022). In addition, in biolinguistics, unlike Labovian sociolinguistics, most contextual and environmental factors are abstracted away. Biolinguists often aim to verify the generative architecture of the language faculty, replete with autonomous syntax and externalised interfaces characteristic of the Minimalist Program (Chomsky 1995b). They hope to do this by means of empirical investigation, using the techniques of neuroscience inter alia.

Corpus linguistics, invigorated by computational tools Saussure could only have dreamed of, homes in on speech corpora and interconnections or collocations between data points. A prominent example of this approach are vector-space distributional models in semantics. Instead of a meaning or semantic value being defined by the association of a word or expression with a discrete set-theoretic object as it is with formal semantics, semantic values are represented as vectors, or, more abstractly, meanings of linguistic items are modelled 'as points or regions in some "meaning space"' (Erk 2020, p. 71). Here, statistical collocational data provides the background for representing semantic structure. A simple case is 'count-based spaces' in which a vector is a sequence of numerical values used to measure the meaning of a word in terms of the words that co-occur with it across multiple contexts. The underlying assumption is that words with the same or similar meanings occur within similar contexts. By representing a word as a matrix of similarity scores within these contexts, we can thus get to its approximate meaning. This perspective breathes new life into Firth's (1957) famous adage that 'you shall know a word by the company it keeps'.

Finally, even in some cases where formalism seems to be doing a lot of the work, such as Optimality Theory (OT), there's a distinctive instrumental flavour to the analysis (Prince & Smolensky 1993). In OT, the basic architecture

Table 1.1. *Basic OT structure*

	Input	Constraint 1	Constraint 2	Constraint 3
	candidate 1	*!		
	candidate 2		*	
☞	candidate 3			*

contains a generator (GEN) that generates an infinite number of outputs or candidates for representation for each input of the grammar. The evaluator component (EVAL) then chooses the optimal output from the set of outputs through a set of ranked, violable constraints, or CON. CON is considered to be universal while the rankings are given by particular languages. In phonology, the inputs are representations from the lexicon and morphology and the outputs are phonetic transcriptions. The basic structure is given in Table 1.1.[20]

In a sense, OT relies on pure combinatorics. It's similar to a mathematical decision procedure. Nevertheless, in spite of the formalism, phonology is very grounded in actual human phonetic possibilities. The candidate list is finite given that phonemes are finite and their combinations limited in real-world languages. The ranked list of constraints is an empirical matter divined from linguistic fieldwork and experimentation.

Even the theoretical linguistic work based entirely on the intuitions of linguists can in part be seen as an empirical exercise, if we follow Devitt's (2006) convincing line of argument that linguists' intuitions aren't Cartesian introspection but rather scientifically honed skills for spotting linguistically relevant structure. The analogy asks us to compare the impressions of a lay person looking over a dig site to that of a paleontologist with years of experience. Naturally, the latter will immediately ('intuitively') see scientifically relevant fragments in ways inaccessible to the novice.

The notion of 'conservativeness' in the nominalist tradition in the philosophy of mathematics is helpful here (Field 1980). The idea is that mathematics aids the sciences in all the ways I mentioned in the previous section on formalisation. More precise models, clearer results, and easier debugging. But the science is *conservative* with relation to mathematics when every claim can be stated without it, in principle. In practice, of course, this makes scientific life very difficult but still possible. More precisely, it suggests that at some high level, the mathematics is dispensable. In other words, all the consequences of the scientific theory (nominalistically stated) derived from the mathematical theory are already consequences of the former without the latter. OT phonology can

20 A star indicates a constraint violation, an exclamation mark after the star means a fatal violation, and a pointing hand an optimal candidate. Sometimes, there's more than one optimal candidate.

most certainly be stated without the tableaux. It's harder to see how this could be for the results of computational corpus linguistics (and part of the reason the original programme was put on hold for a few decades). But again, the statistics are merely meant to reflect the underlying regularities. What's more tricky is the question over the results of FLT and whether the proofs of relative complexity of natural language grammars and constructions are stateable without the formalisms. They're certainly stateable by means of *alternative* formalisms.[21] But the problem is that if mathematical models provide *structural* explanations in empirical sciences and FLT is purely structural, then there might be no remainder for the theory without the maths.[22] We won't settle this question here. In fact, some theoretical linguistics, such as the aforementioned biolinguistics, aims to retain a deep mathematical core of language (Merge or 'narrow syntax') while advocating a naturalist scientific persuasion to the field at large.

Many of the issues discussed in this opening salvo will return within specific contexts in later chapters.

1.5 Cognitive or Social Science?

If linguistics is an empirical science, the next natural question concerns what kind of empirical science it is. The mainstream literature assumes, with Chomsky (1965), that the study of language is the study of (a part of) the mind. Thus, linguistics, on this view, is a cognitive science. Even more, linguistics is often taken to be the study of a particular module of the mind/brain called 'the language faculty'. Of course, this move isn't mandatory. Cognitive linguistics assumes a more domain-general approach to linguistic structure that eschews some of the uniqueness claims of generative theory (Lakoff 1991).[23] Either way, language is considered something internal to the language user. Chomsky infamously stated in *Aspects*:

> Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its (the speech community's) language perfectly and is unaffected by such grammatically irrelevant conditions as memory lim-

[21] For a comparison between generative-enumerative theories of syntax and model-theoretic alternatives, see Pullum (2013).

[22] See Leng (2021) for a recent argument to the effect that mathematical explanations are structural explanations in the sciences generally, and Nefdt (2021, 2022) for a structural realist account of linguistics specifically.

[23] Itkonen (2006) criticises cognitive linguistics as inheriting a fallacy of equivocation from generative linguistics in its definition of 'conventional mental conventions' (Lakoff 1987; Langacker 1991). He argues that cognitive linguistics needs a social normative grounding in order to be successful.

https://doi.org/10.1017/9781009082853.002 Published online by Cambridge University Press

itations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of this language in actual performance. (1965, p. 4)

The idea behind this idealisation was an empirical assumption that many of the mechanisms involved in memory limitations and linguistic dysfluencies of other sorts are independent of those involved in the generation of syntactic structures. What's interesting is that this idealisation not only idealises the internal state of the cogniser but also the state of her community. Essentially, this interpretation of the scientific target of linguistics later developed into an object known as *I-language* in which 'the state attained is a computational (generative) system' (Chomsky 2000, p. 78). The 'I' stands for *internal*, *intensional*, and *individualistic*. The first term simply reiterates the internalism of the field. The second emphasises the need to study grammars qua functions that generate the infinite sets of expressions of language. Lastly, we have the individualistic nature of the scientific object of study. As linguists, we're to be solely interested in the goings-on within the heads of individual speakers and not their environments or relations to others. Part of the motivation for this circumscription is that the alternative, an E(xternal)-language, is allegedly scientifically untenable. Some oft-cited criticisms of this latter possibility involve the politically determined boundaries between languages, poverty of stimulus, and the idea that it might entail an infeasible 'theory of everything' (Chomsky 2000).

Nevertheless, the dual suggestion of a social normative dimension of an empirically based science of linguistics (Sections 1.2 and 1.4) allows for the further possibility of viewing the field as a subfield of social science. There are at least three ways to take this suggestion: (1) incorporating sociolinguistics within theoretical linguistics more firmly, (2) establishing a science of E-language, or (3) exploring the constitutive social normative aspects of language. These options need not be mutually exclusive, of course, even if they're usually packaged that way. Let's explore each briefly.

Sociolinguistics, or option (1), Labovian or otherwise, entails a social ontology but not a normative scientific perspective. Language might be a social object but its properties are identifiable at the social level by statistical, corpus, and/or general social scientific methods. It's a descriptive science. Option (2) takes languages to be external public objects. They're constituted by social conventions. For instance, Dummett holds that 'a speaker has "mastery of a procedure, of a conventional practice"' (1993, p. 69) of a community of speakers when she knows a language. This, to him, is necessary for communication. However, the way to study this social convention need not differ from the tools of standard linguistics or the logical methodology of analytic philosophy of language. Again, without further argument, normativity doesn't enter the picture directly. By contrast, option (3) embraces the normativity.

Itkonen (1978, 1997) motivates the idea that language is a normative social entity and that 'autonomous linguistics' studies it. 'The grammarian does not describe what is said or how it is understood, but what ought to be said or how it ought to be understood. And because the norms (or rules) of language that determine these "ought"-aspects cannot be individual … they must be social' (Itkonen 1997, p. 54). Here, he relies on Wittgenstein's famous argument that the very idea of a private language is incoherent.[24]

In a more contemporary setting, inferentialism adopts a similar normative picture of language and meaning. They too are heavily influenced by Wittgenstein's philosophy. Peregrin (2015), following Brandom (1994), holds that language is constituted by normative inferential rules determined by the 'game of giving and asking for reasons'. Very quickly, as speakers in a community, we're constantly engaging in linguistic (and other) activities that establish 'commitments' and corresponding 'entitlements'. For example, when I make a statement like *Jarda stopped smoking*, I'm automatically committed to a host of other sentences like *Jarda was a smoker*. In terms of methodology, contemporary inferentialism shares features with early proof theory in that it embodies the attempted development of the proof-theoretic approach that characterised the logical constants (in Gentzen-style proof theory, for instance) to also include the nonlogical vocabulary of natural language. This is often referred to as 'strong inferentialism' in that it doesn't restrict the inferential analysis to moves from language to language but also the possibility of inferential rules that govern language–world relationships (following Sellars 1954).[25]

There are issues with both the cognitive scientific and the social scientific interpretations of linguistics. Methodologically, standard linguistics differs from other cognitive sciences and, besides sociolinguistics, from the tools of the social sciences to boot. Ontologically, language has cognitive characteristics as well as a social conventional side. Even Chomsky admits that 'the state of knowledge attained may itself include some kind of reference to the social nature of language' (Chomsky 1986, p. 18). The task is to find a way to pursue both goals without neglecting either. In Chapter 2, I set out one means of achieving a union between these desiderata.

1.6 Outline of the Book

Each forthcoming chapter has at its core a philosophical argument concerning a particular linguistic domain or, in the case of later chapters, the relationship between linguistic domains and other fields. These arguments aren't offered as consensus views or adumbrations of the theoretical questions of the particular topic.

[24] See Chapter 6 for more on Wittgenstein's view of language.
[25] See Nefdt (2018b) for a discussion of the links between proof theory and inferentialism.

In Chapter 2, I'll argue that the central pursuit of theoretical linguistics is the determination of what makes a language a possibly human one. This will take us into the work on recursion, statistical universals, impossible languages, and ultimately bring us to a novel modal metaphysics for linguistic possibility.

Chapter 3 delves into syntactic metatheory. Specifically, I attempt to find a common thread between vastly different contemporary syntactic frameworks from generative grammar and Jackendoff's Parallel Architecture to radical construction grammar and dependency grammar. I'll argue that there's a basic explanatory project in syntax, broader than most accounts but precise enough to offer a genuine point of contact between rival approaches.

The burgeoning field of metasemantics is the focus of Chapter 4. The main pursuit is directed at identifying a viable metasemantic grounding for the field. In this process, I introduce a contextual continuum that connects dynamic and distributional approaches to mainstream formal semantic ones. I also discuss lexical decomposition, underspecification, and supersemantics.

If the syntax–semantics interface looms large in the previous chapter, the semantics–pragmatics distinction occupies Chapter 5. Pragmatics has often been considered 'the waste-basket of semantics'. However, in this chapter, we consider new and old ideas on the semantics–pragmatics interface as well as how these ideas map on to contemporary linguistic pragmatics in the form of relevance theory, OT, game theory, and Bayesian approaches.

The main theme of Chapter 6 is phonology and sign language. The philosophy of phonology is an often-neglected subfield, yet, many of the formal foundations of generative grammar have roots in the methodology of phonology. I'll argue that phonology (broadly construed to include sign, gesture, and haptic modalities) opens up the possibility of thinking of language in terms of a philosophical action theory.

In Chapter 7, we move to a controversial companion to theoretical linguistics. In fact, the central issue of the chapter is whether new approaches in computational linguistics and natural language processing, such as deep learning, have something significant to offer the theoretical study of language. My answer will involve a detour into the philosophy of science, the neglected scientific goal of prediction and its relation to the explanatory project of theoretical linguistics.

The penultimate Chapter 8 will delve into the philosophy of language evolution. Instead of surveying the myriad theoretical options of this topic, we'll challenge the minimalisation trend of the mainstream linguistic approach and argue for a complexity-theoretic interpretation of the scientific target of the field in its stead.

Chapter 9 will conclude with a brief discussion of some topics not covered in depth and an overview of the interconnections between chapters.

Further Reading

Book-length treatments of the philosophy of linguistics are rare. But the following brief list offers readers, new and old, strong introductions to some of the topics covered in this chapter so far.

- A volume by Ruth Kempson, Tim Fernando, and Nicholas Asher (2012), *Philosophy of Linguistics* (Elsevier Ltd), is an excellent place to start with general overviews of philosophical issues lurking within various subfields of linguistics. It contains a number of high-level introductions to issues in the philosophy of phonology, language evolution, computational approaches, syntax, and semantics, written by experts in their respective fields. The coverage is broad, but the depth is strong enough to be accessible to the general reader as well as the seasoned practitioner.
- In terms of formal or mathematical approaches to linguistics, a classic text is still Barbara Partee, Alice Meulen, and Robert Wall's (1993) *Mathematical Methods in Linguistics* (Kluwer Academic Publishers). It covers a large number of tools from within set theory and logic that have been specifically incorporated into linguistic theory in both syntax and semantics. The final set of chapters is especially useful for a deeper understanding of formal language theory, automata theory, and the Chomsky Hierarchy of formal languages.
- For a slightly more recent compendium of mathematical modelling in linguistics, Edward Keenan and Lawrence Moss' (2016) *Mathematical Structures in Language* (CSLI Publications) offers a comprehensive account that includes phonology and a model of Modern Korean. It also covers a lot of the same material as the Partee *et al.* book, such as lattices, set theory, and FLT.
- A more philosophical treatment of the relationship between the formal sciences and linguistics can be found in Marcus Tomalin (2006) *Linguistics and the Formal Sciences* (Cambridge University Press). This book not only serves as a thorough historical overview of the philosophical and mathematical origins of generative grammar but also of its connections to Bloomfieldian formal methods, Hilbert's metamathematical programme, and Goodman's constructional system theory. The evolution of recursion within various epochs of generative grammar is also covered in a careful and precise manner.
- A different take on the relationship between linguistics and philosophy is provided by the recent edited volume by Daniel Altshuler (2022), *Linguistics Meets Philosophy* (Cambridge University Press). This book targets specific phenomena from both philosophical and linguistic angles. The introductory chapter by Barbara Partee is especially illuminating on the history of the interactions between the fields.

- Lastly, Noam Chomsky's (2000) *New Horizons for the Study of Mind and Language* (Cambridge University Press) is perhaps one of the most lucid interactions between philosophy of language and theoretical linguistics in circulation today. In it, Chomsky lays out the philosophical underpinnings of contemporary generative linguistics. He discusses the flaws of traditional philosophical work on language within the analytic tradition and details his alternative naturalistic view of language and linguistics, in which language is considered to be a 'biological object'.