Cambridge University Press & Assessment
978-1-316-51425-2 — The Philosophy of Theoretical Linguistics
Ryan M. Nefdt
Excerpt
More Information

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.

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.

Cambridge University Press & Assessment
978-1-316-51425-2 — The Philosophy of Theoretical Linguistics
Ryan M. Nefdt
Excerpt
More Information

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.

Cambridge University Press & Assessment
978-1-316-51425-2 — The Philosophy of Theoretical Linguistics
Ryan M. Nefdt
Excerpt
More Information

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).

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.

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).

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).