

The Julius Silver, Roslyn S. Silver, and Enid Silver Winslow

DIALOGUES IN ARTS AND SCIENCE

WHAT DO LINGUISTS DO?

Alec Marantz

Silver Professor of Linguistics and Psychology

February 2019



NEW YORK UNIVERSITY

What Do Linguists Do?

What do linguists do?¹ Or, rather, what do linguists do that draws such luminaries as the late Tom Wolfe² to complain so publicly about our enterprise?³ I'll confine my discussion here to what we call "syntax," which is at the core of so-called "Generative Grammar," the way of doing linguistics that has dominated linguistics departments in the US for the past 50 years or so. Linguists deeply engaged in the enterprise of syntax feel that it's obvious what they are doing is not only fascinating, but that it also has led to a set of interesting and important results. However, outsiders can be skeptical, and my own view is that this skepticism can lead to hostility — and bad press — for two reasons. As I will briefly illustrate with some examples from my own research, the actual work of linguists is esoteric and specific. Not everyone gets excited about the details of ant colonies, for example, and not everyone may be kept up at night by the connection between predicate possession in English and adversity causative constructions in Japanese.⁴ However, through their study of esoteric phenomena, linguists claim to be telling us something fundamental about ourselves: about the structure of our minds and brains. Human minds and brains cut closer to home than ant colonies, at least at first blush, so there may be more inherent interest in linguist debunking than persecuting myrmecologists.

Modern syntax finds its roots in work by Noam Chomsky in the 1950's (Chomsky's work itself does not emerge out of the blue but represents a particular intellectual departure from the position of his teacher, Zellig Harris, as well as the work of Nelson Goodman). Many linguists were drawn to the Chomskyan program through his book, *Syntactic Structures*. Within this book one can find places where, if you're destined to become a syntactician, you'll have an "aha" moment, and where, if you're destined to be a skeptic, you'll wonder what's going on. My guess is that the transparency of the generative program to those experiencing the "aha" moment is what leads to an inadequate explication of methodology within linguistics: reading *Syntactic Structures*,⁵ it seems obvious what the game is. But obviousness is, of course, accompanied by a tide of presuppositions and assumptions, so it's worth returning to *Syntactic Structures* to understand what launched syntactic investigations and what escapes contemporary critics of generative grammar.

One can think of the sentences of a language as sequences of units that re-occur. These can be sounds, roughly represented in orthography by letters, or words. With this recognition that sentences can be (partially) described as sequences of recurring units comes the realization that not all sequences of these units are possible sentences in a language. "Understand launched what and generative of grammar," for example, is what linguists call "word salad." How might we describe what rules or generalization separate the strings of words that do represent sentences in a language from those that don't; or, since

it's human speakers of a language that provide the sentences and the judgments about whether the strings could be sentences, how do we describe the knowledge of speakers about the sequences of words that could be sentences? One possibility is that our knowledge is essentially string-based. We know that certain categories of words follow certain other categories — nouns like “cat” follow articles like “the” in English. A mathematically precise formalism for string-centric “grammars” that describe languages in these terms is called a “finite state” grammar. Chomsky asks the question of whether human languages are among the finite state languages, that is, those “generated” (or recognized) by finite state grammars.

The point that Chomsky makes next is what has inspired generations of linguists that followed. He claims we can look at mundane facts about English to show that English, and thus human languages in general, is not finite state (p. 22).

“Let S_1, S_2, S_3, \dots be declarative sentences [not questions] in English. Then we can have such English sentences as:

- (11) (i) If S_1 , then S_2 .
 (ii) Either S_3 , or S_4
 (iii) The man who said that S_5 , is arriving today.

In (11i), we cannot have ‘or’ in place of ‘then’; in (11ii), we cannot have ‘then’ in place of ‘or’; in (11iii), we cannot have ‘are’ instead of ‘is’. In each of these cases there is a dependency between words on opposite sides of the comma.... But between the interdependent words, in each case, we can insert a declarative sentence S_1, S_3, S_5 , and this declarative sentence

may in fact be one of (11i- iii)....”

So, according to Chomsky, English can have sets of *nested* dependencies of the sort, *if...then, either...or, man...is* (as opposed to *men...are*) without a particular limit on the number of such dependencies, and this type of nested dependency cannot be described with finite state grammars.

If you're a kid destined to be a syntactician, you might be immediately intrigued by these facts and fascinated by the phrase structure formalism that can generate these nested dependencies, where phrase structure languages are the next step up in formal complexity from finite state languages. Moreover, we see that what look like unbounded long-distance dependencies from the standpoint of strings end up as *local* dependencies when language operates on phrases, rather than words: although *if* and *then* can be separated by an arbitrarily long sequence of other words, the *if* clause and the *then* clause are essentially sisters within a bigger “mother” constituent, as are the *either* and the *or* clauses. This also applies to the singular or plural subject (*man* or *men*) and the singular or plural verb phrase (with *is* or *are*). Further research in the generative tradition into the 1970's discovered that languages generally traffic in local relations, where the locality becomes apparent once the proper structures are assigned to sentences and the nature of the relations is understood. Moreover, if you're the syntactician to be reading *Syntactic Structures*, you see that you can get started in this field with a pencil and paper (or the 21st century equivalent) and your knowledge of your own native language; the facts you

need to get going are ones you can quickly generate yourself.

Those not destined to become syntacticians might have a different reaction to Chomsky's arguments here, perhaps starting with the observation that *Chomsky actually provides no English sentences*. "S₁" etc. are variables — placeholders for sentences that the reader needs to make up. Or, as I was asked by an astute philosopher in my early years as a linguist, "where's the data?" The worry here is that the mathematical structure of language revealed by nested dependencies in these made up sentences is the mathematical structure of an abstraction (math of math), rather than of the knowledge of speakers of English or of the output of this knowledge. The tradition of syntacticians that followed Chomsky has been to provide actual example sentences in their arguments — the schematic, variable-filled examples show up in mathematical linguistics work, not in journals devoted, e.g., to syntax. And if you're worried about the status of "data" that are generated by the syntactician himself and evaluated against his/her own intuitions, recent work has demonstrated the incredible reliability and replicability of linguistic data, evaluated by the standards of experimental psychology (linguistics is not confronting a replicability crisis of the sort apparent in Social Psychology, for example).⁶

Once we start using sentences as the data of interest for syntax, rather than schemata of sentences, certain objections to generative syntax arise from alternative ideas about what should be done with these sentences. The analytic methods of generative grammar evolve from those of an earlier structuralist

approach to linguistics. On this approach, grounded in anthropology, linguistics is analytic and explanatory, rather than predictive. The goal is to analyze a fixed corpus of sentences from a language by reducing the description of the corpus via rule-based generalizations. For example, recognizing that many English sentences take the form of Name Subject followed by Tensed Verb, followed by Name Object, we could reduce the set of sentences like, “Bill saw John,” “John saw Bill,” “John likes Bill,” ... to a list of names, a list of tensed verbs, and the Subject-Verb-Object “rule.” The rule in some sense explains the presence of the sentences in the corpus.

Rather than consider the rule-based analysis of sentences as a way of systematically reducing a finite corpus, Chomsky pointed to the creative aspect of language use: speakers produce and understand sentences they haven’t heard before. This acknowledgment that any finite corpus of sentences represents an *accidental* sampling of the utterances that speaker can recognize or produce argues instead for a linguistics as part of cognitive psychology, or rather, an enterprise that would engage in prediction beyond corpora of uttered or written sentences. One wouldn’t find sentences like, “If, either the guy who said he’ll complete the job is arriving today, or he’s sending someone else, then we’ll be finished by the end of the week,” in any corpus. However, English speakers can construct as well as comprehend them, and they exhibit the [a [b [c c] b] a] set of nested dependencies that indicate that we have at least phrase structure grammars in our heads (the dependencies: [if [either [the man — is] or] then]).

From prediction of novel data in a given language, generative syntacticians can proceed to predictions about sentences in novel languages. The logic is direct and inescapable: if we can predict how English speakers project from the sentences they hear to novel sentences on the basis of the general structure of language, and if any child from any parents will project English in the same way, then the general structure of language should predict the particulars of all languages. Because language as described here is unique to humans, the general structure of language — which linguists call “universal grammar” — must reflect in some manner what it means to be human, e.g., genetically. Syntacticians thus exploit what they have learned about universal grammar to predict what we should see in all languages: for example, all languages should show the consequences of structures that are not finite state. If a language, then, seems to show a limit on embedding of dependencies of the sort described above, this limit would need to derive from some additional grammatical constraint on structure, rather than from the basic rules of that specific language.

Testing predictions from linguistic theory often has a flavor distinct from what one would expect from an idealized scientific method, in which experiments test competing hypotheses generated from interesting alternative theories (not that any science in practice conforms to this idealization). For example, Yining Nie from my research group was exploring consequences of her approach to causative constructions for French. Causatives, which describe an agent making something happen or bringing something about, strongly

depend on transitivity in her approach, as well as other approaches in the “generative” tradition. The approach predicts that although, “J’ai fait embrasser Jean” (I made/had someone unspecified hug/kiss John), should be fine, “J’ai fait pleurer” (I made someone unspecified cry) should not be, since the resulting structure has no direct object. However, the same analysis would predict that, “J’ai fait pleuvoir” (I made it rain) should pattern with the unacceptable intransitive sentence about crying rather than the perfectly grammatical transitive sentence about hugging. This is incorrect; “J’ai fait pleuvoir” is acceptable, even though pragmatically marked.

Syntacticians call this type of counterexample to a prediction a “puzzle” rather than a disconfirmation or a falsification precisely because there isn’t an interesting alternative theory that made the correct prediction here. Syntacticians live by this type of wrong prediction; puzzles of this sort drive the field. In this case, the puzzle is: what is the right approach to “weather verbs” like “pleuvoir” in French such that the right general account of causatives makes them behave like they have a direct object in causative constructions? If you have linguistic leanings, you’ll be immediately intrigued by this question: do we know things about weather verbs cross-linguistically that might provide a solution to this puzzle? (Indeed, we do.) But if you’re not intrigued, you should at least recognize that there is a puzzle here that could spur investigations of further facts from a variety of languages. Puzzles are counterexamples to predictive linguistic theories that arise in the absence of an alternative theory that predicted them.

An unsympathetic view of linguistic analysis through puzzle-solving has led to a type of critique of linguistics that suggests that these investigations are at best a type of glorified Sudoku and at worst a type of mystical Kabbalah — syntacticians are engaged in an endless search for patterns in language, where the patterns are disconnected from the knowledge that speakers acquire about their languages and exploit in speaking and comprehension. Even a cursory examination of the technical literature undermines this critique: the enterprise is predictive, productive, and competitive. Linguists predict data they don't have, the body of empirical generalizations uncovered by the methodology grows year by year, and alternative accounts of phenomena are in fact pitted against each other, with the losers no longer viable. Progress in linguistics is transparently displayed in our major journals; nevertheless, some scientists and engineers that deal with language still question the legitimacy of the generative linguistic enterprise.

A particular critique of generative grammar comes from computational engineers. If one is interested in processing millions of sentences from the web in the service of predicting the most useful answers to queries on web browsers, for example, one might find that a (compact, useful) finite state description of English covers most of any random sampling of such sentences. For some tasks, moving to a (more complex) phrase structure description provides no noticeable gains, and even models of human reading behavior may not be statistically improved by choosing phrase structure grammars over finite state grammars,

at least over random sets of test sentences.⁷ If you're following closely, you should be thinking, what about sentences like, "If, either the guy who said he'll complete the job is arriving today, or he's sending someone else, then we'll be finished by the end of the week"? Doesn't English involve nested long-distance dependencies of a sort that demand a phrase structure description?

The short answer is, yes, and the computational people that argue with generative grammarians take various approaches to dealing with what Chomsky claims are essential, basic facts about English and about languages in general.⁸ But what should be clear to anyone reading these attacks on linguists is that computationalists are not engaged in the same scientific enterprise as linguists. The linguistic enterprise is about the knowledge of language that underlies everything that a speaker does with his/her language, including not only writing those web pages that serve as data for computational linguistics, but also understanding and making judgments about sentences that are carefully constructed by linguists as test cases to decide between competing theories. For example, suppose our judgments on nested long-distance dependencies stopped at one level of embedding. That might be evidence against phrases structure accounts of language structure, if the judgments were directly related to our knowledge of language, as opposed to being the result of processing overload. Restricting our analyses to narrow subsets of possible data, such as reading time data over random texts or performing a "correct" analysis by a computational parsing program of a large random selection of English texts, doesn't fit with

usual scientific practice. One doesn't know in advance where crucial data might come from, thus, it's necessary to develop interesting competing theories in one's domain of inquiry and design the appropriate experiments to decide between conflicting predictions of the theories. In the case of finite state vs. phrase structure descriptions of English, Chomsky provided the crucial experiment and data in *Syntactic Structures* — the relative paucity of [a [b [c c] b] a] nested long-distance dependencies in any corpora of English hardly speaks to the point.

Another critique of generative grammar arises from an anthropological perspective on language, one that animated the development of linguistics as a field in the US in the first half of the 20th century. Suppose the structure of language is not universal across the human species, nor is it predetermined by our DNA. Suppose instead that the apparent uniformity in structure across well-studied languages reflects the fact that such languages are used by highly literate industrial societies, with phrase structures necessary for the types of functions that such societies employ a language for. Suppose we were to find a people and culture for which the dictates of their society demand the use of a simpler language — not just finite state but actually finite, with a countable number of sentences. That there is such a finite language is the claim of Dan Everett, famously championed by Tom Wolfe.⁹

In a nutshell, Dan Everett claims that the Pirahã of Brazil speak a language that lacks “recursion.” The meaning of “recursion” and its applicability to various structures and proposed rules in linguistics has been a matter

of debate, but in terms of the syntactic structures under discussion in this paper, all would agree that sentences of arbitrary form (S_1 in *Syntactic Structures*, e.g.) embedded within sentences of arbitrary form would constitute a type of recursion. Everett denies that Pirahã grammar generates such structures, nor do Pirahã sentences contain noun phrases (of arbitrary structure) embedded within noun phrases, or adjective phrases within adjective phrases. In fact, given the constraints Everett proposes for Pirahã, the language is essentially finite: one could list all the sentences of the language.

Although trained as a generative linguist (and although the author of excellent descriptions of Pirahã written within the Chomskyan generative tradition, in which Pirahã displays a recursive grammar of the usual sort), Everett writes his current description of the Pirahã as an anthropologist. He tries to explain the properties he observes in the Pirahã language by referring to aspects of Pirahã culture. In particular, he attributes to the Pirahã a commitment to the here and now, constituting a general cultural constraint against talking or thinking about times or places in the distance, specifically those outside one's own experience and the experience of the current set of tribe members. In his writing, Everett tries to connect this general principle to various additional aspects of the Pirahã language and culture — lack of color terms, lack of numbers, etc.

Although I am not an anthropologist, I can say that Everett has written mediocre to poor anthropology. I can say this because it's the cultural anthro-

polologist's job to provide the reader an understanding of the culture being described — to tell a compelling story that allows us to understand the difference between another culture and our own — and I don't get any feeling for the Pirahã or understanding of their culture from reading Everett. For the present discussion, the largest failure in Everett's anthropological explanations is his notion that cultural constraints can explain behavior strictly adhering to the constraints when the possibility of a violation of the constraints is apparent. I would buy the cultural constraint explanation better if Everett included stories about how the tribe dealt with transgression and rebellion. A story about how the group dealt with, say, a rebellious teenager who code-switched between Pirahã and Portuguese (or who yelled out “large happy fish” in Pirahã, with recursive noun phrase structured (see below)) would go a long way convincing me that Everett has described an alternative way of being human, rather than an exotic, inscrutable people beyond our understanding.

But as I claimed earlier, even if a poor anthropologist, Everett is a good linguist. As such, he knew he needed to provide a (semi)formal account of the grammar of Pirahã – how to describe both the structure of Pirahã as well as the limited nature of its syntax. His analysis of Pirahã actually involves claiming Pirahã is just like every other language, except that it has a version of a mechanism that other languages use that, in Pirahã, limits the level of embedding of words within phrases. Essentially, Everett claims that every noun, verb, or adjective in Pirahã needs to be “licensed” in its syntactic position by being in

construction with (structurally close enough to) an “evidential” marker — a morpheme that attributes knowledge of the propositional content of a sentence to the speaker or to someone else (like “reportedly” in English). There’s essentially one evidential marker per sentence, at the level of the whole sentence, so any noun, verb, or adjective must be high enough within the constituents of the whole sentence to be visible for licensing. The meaning of “restrictive” modification of a noun involves a computation that is semantically recursive (basic modification is “intersective”: a “green tree” refers to a member of the intersection of the things that are green with the things that are trees). If a language allows modification of a noun, then the result of such modification, being of the same semantic type as a noun, may serve as the input to additional modification. In Pirahã, however, if there’s a single evidential marker on the verb of a sentence (and one verb per sentence), then recursive modification puts elements in the phrase containing the modified noun deeper and deeper within this phrase and farther and farther from the evidential. In a structure with a verb + evidential and a noun phrase, say, an object, a [large [happy [fish]]] would include at least one adjective or noun too far from the evidential for licensing, even though the meaning of such recursive modification can be computed by the Pirahã, Everett claims. The requirement for licensing by an evidential, then, limits the degree of modification to one — no recursion allowed.

In short, Everett’s own formal account of Pirahã syntax makes Pirahã conform to “generative” claims about universal grammar. The limits on recur-

sion and on the “creativity” of Pirahã speakers with respect to the creation of new sentences follow not from a direct constraint on the grammar that would make Pirahã finite-state in its essence, but rather from a particular form of licensing constraint that acts to restrict recursive structure building. However, that is not to say that most linguists would agree with Everett’s account.¹⁰ In addition to pushing an anthropological stance in his analysis of the Pirahã language, Everett has denied the applicability of standard linguistic methodology for the study of Pirahã. Standard informant work on languages involves asking for judgments of acceptability, for example, of novel examples constructed by the linguist. Everett proposes that such methodology would be inappropriate for the Pirahã. Rather, he seems to suggest that we return to the structuralist methodology of working with corpora of utterances from speakers, which is a bit like waiting for sentences like, “If, either the guy who said he’ll complete the job is arriving today, or he’s sending someone else, then we’ll be finished by the end of the week,” to emerge from the mouths of English speakers. That is, Everett refuses to test competing theories about the structure of Pirahã using the standard scientific methodology of linguistics.

The project envisioned in *Syntactic Structures*, that of characterizing human cognition by exploring the mathematical properties of human language, continues today, primarily among computational linguists.¹¹ Although Chomsky convincingly demonstrated that natural language requires more computational “power” than that of context-free phrase structure grammars,

there are apparently interesting differences among levels of complexity in the range between context free grammars and the power of Turing Machines, and the exact place of our linguistic competence in the hierarchy is still unknown or unproven. However, this question does not drive my own work in Universal Grammar, which is more concerned with the ontology of elements that are combined to make words and phrases, and in the specifics of the interactions among these elements, as opposed to the general mathematical characteristics of these interactions. That makes me somewhat of an “empiricist,” as Chomsky informed me when I was his student, whether in a neutral or evaluative tone I’m still not sure to this day.

At the moment (Spring semester of 2019), I’m investigating the nature of “locality” constraints on the interactions of the combinatory elements of language: the “morphemes.” This investigation is directly connected to Chomsky’s observation from *Syntactic Structures* I explained above, that the apparent long-distance relation between “if” and “then” in a conditional sentence is actually a very local relation between the “if” clause and the “then” clause, once the hierarchical structure of language is acknowledged. Certain types of interactions between morphemes require a very local relationship, similar to the connection between “if” and “then.” For example, some suffixes restrict the meanings of the stems to which they attach. Although “novel” can refer both to the content of a book (“this novel is very interesting”) or to the book as an object (“this novel weighs half a pound”), the verbalizing suffix *-ize*

chooses the “content” reading. One can “novelize” someone’s life, but one can’t “novelize” these sheets of paper into a physical book. If we add another suffix to nominalize the verb, to create “novelization,” the final suffix isn’t local enough to the stem, “novel,” to choose the physical object reading, which is thus no longer available to further suffixation after the addition of *-ize*. We can’t talk about the “novelization” of reams of paper.

These sorts of observations yield interesting insight into a variety of puzzling phenomena. For example, the prefix *re-* seems to mean the same thing as “again”: “I re-opened the door,” seems equivalent to, “I opened the door again.” However, *re-* only expresses one of the possible meanings of “again,” that of restoring the previous state of an object, not that of repeating an action. So, “I re-opened the door,” means, “I opened the door, and it was open before.” As a result, *re-* does not attach to verbs expressing activities, rather than changes of state. Although I can say, “I danced again last night,” I can’t say “I re-danced last night.”

Now the careful reader should be thinking, isn’t “dance” in “re-dance” close enough to *re-* to allow it to choose a special “do again” meaning for *re-*? *Re-* is right next to *dance*, after all.

That is, couldn’t I make “re-dance” mean “dance again” if I wanted to, and get other speakers to agree? The answer is actually, no, you can’t. The hierarchical arrangement of morphemes in language is called the “syntax” of the

language. The same hierarchical syntactic structure is interpreted (as Chomsky says, “externalized”) in sound/writing/sign on the one hand and in meaning on the other. So, we can find clues as to the arrangement of morphemes relevant for meaning in the way that morphemes are pronounced — both the meaning and the pronunciation are computed from the same structure in highly constrained ways. The “restitutive” (indicating restored state) *re-* prefix in English bears stress (compare unstressed *re-* in “I returned the bottle” with stressed *re-* in “I re-turned the knob to the open position”). In English, the stress on a prefix indicates that it belongs to a separate syntactic domain from the stem to which it attaches; a full analysis of *re-* would show it belongs to the syntactic domain of the direct object independent of the verb. Belonging to separate syntactic domains, as indicated by their pronunciation, *re-* and *dance* cannot influence each other semantically as they could if they belonged to the same domain. Even though they look like they’re next to each other, like “if” and “then,” *re-* and *dance* are in fact not syntactically local in the relevant sense for semantic interpretation, if *dance* were to want to choose the repetitive rather than the restitutive reading of *re-*.

Investigations into the properties of English words hold strong implications for the architecture of grammar and the nature of language. However, the sorts of locality constraints at issue here on the interactions of morphemes do not immediately hold obvious implications for the nature of human cognition, unlike Chomsky’s conclusions that human languages involve more computa-

tional power than context free grammars. My type of research is more standard cognitive science: investigating the representation and use of language without explicitly making claims that the nature of linguistic representations and computations are either unique to humans or unique to language. Or, I should say, my type of research is standard cognitive neuroscience, since these investigations include testing hypotheses via connections to neural mechanisms. For example, one implication of the theory of morphemes and their combination that I've been developing on the basis of the sort of evidence about *novelization* and *re-dance* that I sketched above is that English speakers must analyze words like "excursion" into stem and suffix morphemes "excuse" and "ion" even though the putative stem "excuse" occurs only in the word "excursion." My lab has confirmed this, as well as other surprising predictions about the neural processing of morphologically complex words, using a brain monitoring technique called "magnetoencephalography," which measures the magnetic fields generating by electrical activity in neurons in the cerebral cortex.¹² Nothing about the linguistic enterprise demands that one marshal evidence from brain signals to support one hypothesis over another. But, more crucially, nothing about linguistics tells us that we won't find crucial evidence for a theory in brain responses either.

What do linguists do? We're not engineers and we're not anthropologists; critiques of linguistics based on the goals of these disciplines are therefore at best baffling. We are cognitive neuroscientists and, perhaps with vision

scientists, the pioneers of this discipline. We operate at many levels of abstraction, from the quite general claims about human cognition based on computational linguistic analysis of the mathematics of human grammars to specific hypotheses about the neural responses to “excursion.” Like the findings of any science, the observations of linguists from the behavior of “if...then” sentences to the ungrammaticality of “re-dance” should appeal to our innate curiosity as humans and to our desire to understand more about the world around us. Chomsky suggests, following the rationalist tradition of Descartes among others, that linguistics may hold special appeal to the extent that language represents a species-specific, genetically encoded capacity; understanding the nature of language is part and parcel of an appreciation of what makes us human. However, I’d like to believe that Chomsky’s discussion of nested dependencies in English syntax, or my more humble description of the (im)possible meanings of *novelization* and *re-dance*, would be sufficient to evoke the foundations of a fascinating and productive field of study.

1 Thanks to Wayne O’Neil, Jay Keyser, and Noam Chomsky for comments on an earlier version of this paper, and to Ellie Abrams for editorial assistance.

2 Wolfe, Tom. *The kingdom of speech*. Random House, 2016.

3 Of course, Wolfe claims that both Darwin and Chomsky are charlatans in his book, putting linguistics in the same company as evolutionary biology.

4 Wood, Jim, and Alec Marantz. “The interpretation of external arguments.” *The verbal domain* (2017): 255-278.

5 Chomsky, Noam. *Syntactic Structures*. Walter de Gruyter, 2002. [Original edition published in 1957.]

6 Sprouse, Jon, Carson T. Schütze, and Diogo Almeida. “A comparison of informal and formal acceptability judgments using a random sample from Linguistic Inquiry 2001–2010.” *Lingua* 134 (2013): 219-248, Sprouse, Jon, and Diogo Almeida. “Assessing the reliability of textbook data in syntax: Adger’s Core Syntax.” *Journal of Linguistics* 48.3 (2012): 609-652.

7 Frank, Stefan L., and Rens Bod. “Insensitivity of the human sentence-processing system to hierarchical structure.” *Psychological science* 22.6 (2011): 829-834. But see, e.g., Roark, Brian, et al. “Deriving lexical and syntactic expectation-based measures for psycholinguistic modeling via incremental top-down parsing.” *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1*. Association for Computational Linguistics, 2009.

8 E.g., Frank, Stefan L., Rens Bod, and Morten H. Christiansen. “How hierarchical is language use?.” *Proceedings of the Royal Society B: Biological Sciences* 279.1747 (2012): 4522-4531.

9 E.g., Everett, Daniel, et al. “Cultural constraints on grammar and cognition in Pirahã: Another look at the design features of human language.” *Current Anthropology* 46.4 (2005): 621-646.

10 See, e.g., Nevins, Andrew, David Pesetsky, and Cilene Rodrigues. “Pirahã

exceptionality: A reassessment.” *Language* (2009): 355-404.

11 See, e.g., Graf, Thomas. “Grammar Size and Quantitative Restrictions on Movement.” *Proceedings of the Society for Computation in Linguistics (SCiL) 2018* (2018): 23-33.

12 Gwilliams, L., and A. Marantz. “Morphological representations are extrapolated from morpho-syntactic rules.” *Neuropsychologia* 114 (2018): 77-87.