

1 The Origins of Phrase Structure Constituency

In a remarkably prescient paper written for the 1989 Tilburg Conference on Discontinuous Constituency, David Dowty makes the following observations:

No assumption is more fundamental in the theory (and practice) of syntax than that natural languages should always be described in terms of constituent structure, at least wherever possible. To be sure, certain kinds of cases are well-known where constituent structure of the familiar sort runs into problems . . . But even here, the strategy has usually been to accommodate these matters while modifying the received notion of constituent structure as little as possible.

There are two things that worry me about the situation syntactic theory finds itself in. Since hierarchical syntactic structure is so often assumed, syntacticians don't usually ask questions—at least beyond the elementary syntax course—as to what the nature of evidence for a constituent structure in a particular sentence in a particular language is: we just take whatever structure our favorite syntactic theory would predict as the expected one for the string of words in questions—by the current X-bar theory, etc.—unless and until that assumption is contradicted by some particular fact.

My second concern is closely related: I suspect syntacticians today have almost come to think of the primary empirical data of syntactic research as phrase structure trees, so firm are our convictions as to what the right S-structure tree for most any given sentence is. But speakers of natural languages do not speak trees, nor do they write trees on paper when they communicate. The primary data for syntax are of course only *strings* of words, and everything in syntactic description beyond that is part of a theory, invented by a linguist. (Dowty 1996b, 11–12)

Dowty's misgivings about the default status of phrase markers as the appropriate data structure for representation of the fundamental syntactic properties of sentences are, as we argue in this book, entirely justified, and in succeeding chapters we present a considerable body of additional phenomena, beyond those he cited in his paper, to support this negative view of configurational representations. But it is worth beginning this investigation by addressing the question of *why* we have this state of affairs in the first place. Where did the notion of phrase structure (PS) constituency come from, and why has it maintained such a grip on the analytic imagination of syntacticians—to an extent which makes it seem reasonable to posit all manner of complex, stipulative operations just to

accommodate phenomena which stubbornly resist description in terms of hierarchical trees? As Dowty notes, various “fixes” have been proposed, essentially as add-ons to a phrase structure basis—indeed, in all but the most recent phase of transformational grammar, this was the essential architecture of the combinatoric component—but it still seems to be regarded as perverse, in some foundational way, to reject hierarchical phrase structure as the model of syntactic representation.

In this introductory chapter, we sketch what seems to us to be the main line of development in generative linguistics (and its precursor) which led to this state of affairs. Our view is that the notion of phrase structure hierarchy as a theoretical construct originates as a reification of analytic practice in the American Structuralist tradition originating with Bloomfield and his students. In the work of the Structuralists, this practice had no special theoretical status but rather represented a method of description which extended procedures applied and honed in the domain of morphology to the larger units beyond the word that the Structuralist found it necessary to recognize. And these morphologically based practices themselves originated, ultimately, in the linguistic relativism that Bloomfield’s teacher, Franz Boas, introduced into the history of the field in the work he carried out and sponsored on the native languages of the New World. Our point in the following discussion is that phrase markers evolved out of a set of descriptive practices which reflect the analytic methods of Boas’s intellectual descendants, rather than from some strongly motivated theoretical foundation. The long-established consensus that hierarchical constituency is the optimal data structure for syntactic representation is, in our view—and as per Dowty’s caveat—a largely contingent development that calls for rethinking at a fundamental level.

1.1 The Boasian Era

Without any question, the dominant figure in American linguistics in the early twentieth century was the celebrated cultural anthropologist Franz Boas. When Boas first began working on native languages of North America, the dominant view of linguistic form was heavily Eurocentric, reflecting the views of the major nineteenth-century comparativists. On this view, the classical languages (Greek, Latin) and certain favored modern languages (e.g., German) embodied the “essence” of linguistic form, and all languages were to be described using, for example, Latin paradigms as models.

Boas took a radically different approach, concisely formulated in his remark in the introduction to the massive collection of descriptive grammars, *Handbook of American Indian Languages*:

[T]he method of treatment has been throughout an analytical one. No attempt has been made to compare the forms of the Indian grammars with the grammars of English, Latin, or even among themselves; but in each case the psychological groupings which are given depend upon the inner form of each language. (Boas 1911, 68)

By “psychological groupings,” Boas meant the grammatical categories revealed by morphological analysis, the point being that the authors of each of the grammatical descriptions in the *Handbook* understood the language so described to have its own formal organizing principles, with no language regarded as superior to, or entitled to serve as a model for, the analysis of any others.

The challenge in taking such an approach is of course to settle upon a method of analysis which does not make any reference whatever to some predetermined checklist of categories. From this point of view, one could view the movement which dominated American linguistics in the 1930s through the 1950s, *American Structuralism*, led by Leonard Bloomfield, one of Boas’s two most important students, as in effect the operationalization of Boas’s notion of “the inner form of each language.” This fleshing out of Boas’s conception could be summarized as the principle that any language was to be described in terms of classes of forms with parallel distributional properties. Thus, word parts, words, and strings of words could be assigned to equivalence classes based strictly on where they could appear in words (morphology) and sentences (syntax). The rules governing the arrangement of members of these equivalence classes (or “form classes”) in any language constituted the grammatical rules of that particular language.

In structuralist analysis, morphological analysis—as expounded in detail in Nida’s monumental 1949 text—was the model for syntactic analysis, a fact explicitly acknowledged in, for example, Zellig Harris’s seminal “From Morpheme to Utterance” (Harris 1946) and Rulon Wells’s more influential 1947 paper (Wells 1947) on the descriptive device he referred to as *Immediate Constituency* (IC). Harris (1946, 161) begins his paper with the following remarks:

[T]his paper presents a formalized procedure for describing utterances directly in terms of sequences of morphemes rather than single morphemes . . . At present, morpheme classes are formed by placing in one class all morphemes substitutable for each other in utterances, as *man* replaces *child* in *The child disappeared*. The procedure outlined below consists, essentially in extending the technique of substitution from single morphemes (e.g. *man*) to sequences of morphemes (e.g. *intense young man*) . . . When applied to a particular language, the procedure yields a compact statement of what sequences of morphemes occur in the language, i.e., a formula for each utterance (sentence) structure in the language . . . The reason for a procedure of the type offered here is not far to seek. One of the chief objectives of syntactic analysis is a compact description of the structure of utterances in the given language.

Here and elsewhere in his paper, Harris is quite explicit that the “method of substitution” he alludes to as an extension of morphological analysis has as its objective a concise—that is, maximally general—statement of the distribution of words and word sequences in whatever language it is applied to. No theoretical content is claimed for the distributional statements arrived at in this fashion; Harris’s goal, as he makes clear

throughout the discussion, is the explicit codification of standard Structuralist analytic practice. As Harris (1946, 161) observes, “the proposed method does not involve new operations of analysis. It merely reduces to writing *the techniques of substitution which every linguist uses as he works over his material*” (emphasis added).

Essentially exactly the same method is presented in Wells (1947), more widely cited perhaps as a result of his possibly more intuitive use of bracket notation rather than the seemingly mathematical formulæ employed by Harris. So far as content is concerned, Wells’s and Harris’s respective papers share largely the same central premise: what Wells (1947, 82) identifies as “the simple but significant fact on which we base our whole theory of I[mmediate] C[onstituent]s . . . : that a sequence belonging to one sequence-class A is often substitutable for a sequence belonging to an entirely different sequence class B.” By “sequence class,” Wells is referring somewhat neutrally to sequences of morphemes, but in the context of the specific data and examples he provides, a sequence class is essentially a sequence of words with a parallel distribution with a sequence of words that has the same distribution as a single word taken to be an “expansion” of the sequence class to which that single word belongs. What Wells is in fact setting up is a definition of constituency in which constituents of a sentence are determined by parsing the sentences into strings whose distribution parallels that of single words in a range of contexts.

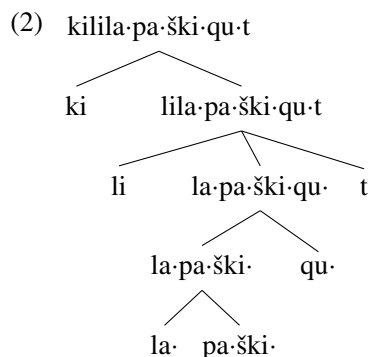
We can, however, be a bit more specific about the conceptual origins of constituency in Structuralists’s standard practice. A substantial number of major Structuralists—including theoretical leaders such as Bloomfield and Hockett, and influential figures including Eugene Nida and Kenneth Pike—had continued the Boasian tradition of field linguistics with native languages of the western hemisphere, such as those of the Algonkian, Iroquoian, and Athapaskan families. In these languages, root morphemes were preceded and/or followed by sequences of affixes, many arranged in paradigmatic classes (“position classes”). Faced with such complexity, the standard methods of descriptive analysis pioneered by (post-)Boasian field linguists involved the elicitation of both inflectional and what might be called derivational paradigms, identifying recurrent pairings of phonetic form and grammatical function, where the analysis proceeds from the outside inwards peeling off layers of affixation until one arrives at the “roots” and “stems” referred to as such in, for example, Bloomfield (1933) and Nida (1949). Although the conclusion of this analytic practice was typically a statement of how morphological composition may proceed in the construction of lexical items, the practice itself orients the analyst to approach lexical items as a cluster of parts which are to be separated out in a series of steps, where at each step some of the material has been segmented off, a remainder which is then to undergo further analysis. This step-by-step method resolving words into sequences of morphemes is difficult to apply in any other way, for as a matter of practical necessity, the alternative would be

to try to carry out segmentation medially, with unidentified phonological material on both sides of any given internal string of phonological segments, greatly increasing the uncertainty of the results.

The result of this kind of analytic practice ultimately had far-reaching implications for syntax. As Joos (1957, 185) notes, the Structuralists did not really start approaching syntactic description seriously until they had acquired considerable experience with morphology, where their distributional methods beyond the level of phonological segments had been honed—a fact which clearly predisposed them to use extensions of the familiar methods they had, in so many cases, brought to bear on morphologically very complex field languages and to view syntax from the perspective of those methods. Thus, Nida (1949, 86), for example, observes that “[t]he distribution of any morpheme must be given in terms of its environment, but some of its environment may be important and the rest unimportant. *This is true of both morphology and syntax.*”¹ In particular, the standard technique employed in descriptive linguistic practice when one has to analyze a word in a language different from one’s own native language(s) has typically been to work from the outside in, observing what happens to word meanings when partials at the beginning or the end of a word are replaced while the rest is held constant, or, conversely, what constant contribution to meaning corresponds to a given stretch of phonological material at the edges of words across a number of data points. Once one has separated out the material on the periphery, the analytic procedure can be iterated, until the form is completely parsed.

This “decompositional” preference, arising from the practical methodology adopted by grammarians guided by distributional criteria, would have made it extremely natural for linguistic objects to be seen as assembled out of smaller parts, which in turn might themselves comprise subparts. Nida (1949, 101) in fact presents just such a treelike analysis for the Totonac word *kilila-pa-ški-qu-t* (‘my necessity of loving them reciprocally’):

1. None of this is to say that Structuralist methodology was the sole contributor to the notion of hierarchy in human languages; thus, Percival (2018) identifies, as an important strand of Bloomfield’s thinking, the influence of the psychologist and philosopher Wilhelm Wundt in his major work *Völkerpsychologie: Eine Untersuchung der Entwicklungsgesetze von Sprache, Mythos und Sitte* (Wundt 1900), arguing that Bloomfield’s earlier major work on linguistics, published in 1914, was heavily indebted to Wundt’s insistence that sentences have a primarily binary internal structure independent of the logical content the sentence expresses. In particular, as Percival notes, Wundt’s own thinking was heavily influenced by the structure of formulae in symbolic logic, whose connectives are primarily binary, and in which propositions are typically decomposed into a predicate and its argument. But it is unlikely that Wundt’s logic-based binary view of the analysis of sentences and its influence on Bloomfield was the primary driver of the hierarchical analytic picture of sentence structure held by the leading Structuralist theorists. Percival himself acknowledges, in connection with the Structuralists’s “preference for binary analysis” (clearly evident, for example, in Wells’s formulation of IC analysis, to which we turn directly), that morphological practice played a major role in this attitude.



From the analytic perspective given in the work by Harris and Nida, the graphic object in (2) is nothing more or less than a compact statement of the outer-to-inner segmentation of the complex Totonac word as motivated by the distribution of the substrings separated out in (2) on the basis of the recurrent semantic contributions of those subparts across the whole recorded Totonac lexicon: there is a “nucleus” comprising the immediate constituents *la·*, the reciprocal prefix and the verb stem *pa·ški·*, ‘to love’; this sequence was in turn an immediate constituent, with the inflectional object suffix *qu·*, of the verb *la·pa·ški·qu·*, ‘reciprocate love of someone’; and so on. These partials are identified via the distribution of form/meaning pairing, which was the major analytic tool in the Structuralists’s arsenal, whereby a phonological sequence *ki-* has the constant meaning in all its occurrences of first-person possessive, *li- . . .t* can be separated out from the rest of the word it appears in as a derivational circumfix denoting the necessity of the nominal it derives by combination with some verb, and so on. Essentially, the immediate constituent tree in (2) is nothing more or less than a graphic summary of this distribution of phonological/semantic covariation based on extensive elicitation and inspection of lexical and textual material.

Analyses of the sort summarized in (2) led to a certain view of word formation summarized in the following observation by Nida (1949, 98):

[L]anguages with an extensive morphological structure frequently exhibit well-defined structural layers . . . The principle division is between derivational and inflectional formations. The derivational formations may in turn be divided between formations of the nuclei (these are compounds) and constructions consisting of nuclei and nonnuclei.

The history of the analyst’s decomposition of the word into its distributionally determined subparts was thus seen as providing, in reverse order as it were, a guide to the morphological steps by which the word had been built up on the basis of its atomic root elements and affixes.

In view of the foregoing, it is hardly surprising that Structuralists’s efforts to extend their methods to description and classification at the sentence level took the same form as their morphological descriptions. The basic thrust is evident from Wells’s exam-

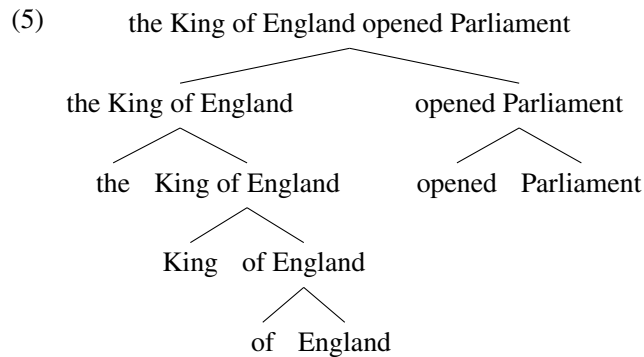
ples. Wells's central illustration of the method of immediate constituent analysis, for instance, starts with the sentence in (3):

(3) The King of England opened Parliament.

and breaks it down into immediate constituents based on substitution possibilities along the lines indicated. As Wells formulated his results, they took the explicit form of the following unlabeled bracketing (Wells 1947, 84):²

(4) [[the] [[king] [[of] [England]]]] [[opened] [Parliament]]

(4) simultaneously documents the *defensible* breakdown steps to arrive at the terminal yield of the analysis and displays the token existence of the strings derived in the breakdown at every level of the analysis. What is important in terms of the future development of linguistic theory in the post-Structuralist phase is that such objects essentially identify a relationship of “containment” among strings—simultaneously identifying, in particular, which strings any given word sequence is part of and what substrings it contains itself. These containment relations correspond at the syntactic level to exactly the information represented in the morphological analysis summarized in (2), and can be presented in the same graphical format.³



Such a labeled bracketing, by its very nature, depicts both the largest substring constituent parts of the sentence and the largest substring constituents of each of those parts, and so on. Because one is working top-down, so to speak, one first reveals the yield of the initial breakdown steps, then the breakdown of the latter into the largest strings they contain that meet Wells's substitutability criterion, repeating the procedure

2. Wells uses the vertical bar | instead of left and right brackets, but the information represented is the same.

3. Indeed, Nida (1949, 84) presents just this kind of tree in his representation of the structure of the sentence *Peasants throughout China work very hard*.

down to the individual words.⁴ And because the critical relationship illustrated is containment, by far the most natural way to depict the totality of breakdown steps is by bracket annotation of the word sequence constituting the sentence.

IC analysis of the sort typified in (5) is clearly a major source of the dominant role that the notion of phrase structural constituency plays in modern linguistics, although during the Structuralist era it had, as already noted, no particularly rigorous methodological foundations. Wells's own characterization of the criteria by which alternative constituent analyses are arrived at is not particularly systematic and never gets around to explaining why such criteria are salient in the first place. At no point in his paper does Wells offer anything like a theoretical frame within which his classification criteria for determining the immediate constituents can be taken to yield a *better* result than some other set of criteria.⁵ It is all the more interesting, then, that this exclusively methodological view of constituency was essentially enshrined as the canonical language of syntactic representation in the generative paradigm which succeeded Structuralism—a paradigm whose defining literature repeatedly emphasizes its theoretical content, in contrast to its taxonomic predecessor.⁶

1.2 Transformational Grammar: The Structuralist Legacy

It is by now fairly well accepted that many of the supposed innovations in transformational grammar were taken over, often with only cosmetic modification, from earlier

4. Wells does take the position that IC analysis can be construed as either a bottom-up or top-down process. But his examples uniformly reflect a top-down series of analyses, and when he says that “[t]his is the fundamental aim of IC-analysis: to analyze each utterance and each constituent into maximally independent sequences—sequences which, consistently preserving the same meaning, fit in the greatest number of environments” (Wells 1947, 88), it seems clear that the linguist is expected to identify a set of data and then apply the method illustrated in Wells's treatment of (3). Indeed, it is difficult to see how the particular methods Wells employs could be carried out in any *but* a top-down manner.

5. For example, the standard metric for evaluating scientific hypotheses—predictive success over the range of data to which the hypothesis applies—is not in play here, because there is no predictive content to Wells's proposal, nor, it is apparent, was there intended to be any. What Wells was outlining, as he makes clear throughout the paper, was a program for systematically applying the substitution criterion for defining constituents illustrated in some detail in Harris (1946). As such, it reflected the Structuralists's essentially exclusive concern with explicit procedures for classifying data.

6. Hyemes and Fought (1981, 125) make the rather odd claim that Wells was consciously engaged in linguistic theorizing because he uses the word *theory* in the first paragraph of his 1947 paper. It is instructive to recall what Wells actually says in that paragraph:

We aim in this paper to replace by a unified, systematic theory the heterogeneous and incomplete *methods* hitherto offered for determining IMMEDIATE CONSTITUENTS . . . The unifying basis is furnished by the famous concept of patterning, applied repeatedly and in diverse special forms. (Wells 1947, 81; emphasis added)

It is quite evident that by *theory* here, Wells intends nothing more than a systematic procedure, based on the parallel patterning of certain expressions and other much shorter ones, viz., that both can occur in a sufficiently large and diverse range of grammatical environments.

work in Structuralist linguistics on the one hand and mathematical logic on the other. Pullum (2011), for example, notes that Chomsky's watershed publication, *Syntactic Structures* (Chomsky 1957), offers a rule system which is, in the end, a more transparent version of formal generative technology than had been first devised in the mid-1940s by the logician Emil Post (see Post 1943, 1947). But what is particularly striking about the system which Chomsky presented in *Syntactic Structures* is the degree to which it took over what had been the methodological toolkit of Structuralist linguistics as the underlying component of syntactic theory. This is somewhat ironic, since generative grammar is standardly understood as having sharply broken with the Structuralists' practice of defining "discovery procedures" (of the sort described above) by emphasizing the primacy of "predictive success" as the criterion which, as in other nomological-deductive enterprises, grammatical analyses are to be evaluated against.

It is worth noting in this connection that early in generative grammar's career, after so-called generalized transformations were eliminated from the framework, the underlying context-free rule system (the so-called rules of the base) was *entirely* responsible for the infinite cardinality of the set of sentences each human language was taken to comprise. The open-endedness of the string sets that a language comprises is a by-product of the availability of chains of PS rule applications which expand some phrasal category *X* to yield, at some point in the series of rules applications, a string containing *X*. And collectively, these base rules were essentially a formalization of Harris's and Wells's immediate constituent analysis via a reader-friendly renoting of Post's generative technology (as per Pullum 2011).

In this context, it is useful to consider just what was involved in the transition from Wells's unlabeled bracketing representations to phrase markers with category-labeled nodes which became common at the end of the 1950s. The notion of constituent structure that Chomsky had inherited from the Structuralists reflected the containment relationship displayed in the hierarchically organized decomposition of sequences, as per (2) and (5). But the replacement of Wells's unlabeled bracketing capturing containment relationships among these constituents by labeled brackets made it possible to link such structures to the virtual automata which were increasingly coming into use as mathematical models of the generative systems enumerating those structures, where the categories labeling the brackets could be taken to correspond to the states of the automaton, and the PS rules to analogues of the transition functions defining such automata. Still more importantly, the category labels attached to bracketed strings, corresponding to the nodes in phrase markers, enable one to state generalizations over the behavior of all strings of words so labeled.

This point deserves some amplification, since it bears on the way in which a data structure made available in one analytic paradigm both fed into and helped shape the thinking of researchers working in succeeding paradigms. Phrase structure trees,

widely adopted as a notation for syntactic representation soon after the appearance of *Syntactic Structures*, convey the same information content as labeled bracketings—but do so via a graphic organization of information which renders transparent the “genealogical relationship” among the category-labeled nodes in a way that bracketed objects do not. It is far easier to identify a string of words as descendants of a single category label when tree notation is used than with a flat string annotated with multiple subclusters of delimiters—and this, of course, matched perfectly with the appeal of transformational rules, originally proposed by Zellig Harris (1951), that Chomsky renotedated and presented as the central innovation of his own framework. As structure-dependent mappings between hierarchical constituent representations, transformations depended, from their earliest formulations, on the possibility of identifying a string of words as the exhaustive descendants of a given category node.⁷

A key aspect of the role played by constituency in the transformationalist era is the way in which transformations, though nominally the major theoretical attraction in the new paradigm, frequently wound up playing the role of diagnostic probes for underlying phrase structure. Early textbooks, such as Jacobs and Rosenbaum (1968) and Akmajian and Heny (1975), argue for constituent structure at various points by pointing out various examples which appeal to movement in their derivations. The fact that transformations were taken to apply to major categories meant that the strings of words affected by such transformations were ipso facto analyzable as tokens of one or another of those categories. The role of the domination relation in identifying word sequences as analyzable with respect to some higher category node—and hence potentially part of the structural description of one or another transformational rule—ensured in turn that transformations could be used as diagnostic probes for phrase structure. But again, there was actually nothing fundamentally new in any of this. The Structuralists had, after all, insisted that shared distributional properties were the essential basis for taking (sequences of) grammatical formatives to belong to the same “form classes,” and in Wells’s analysis this diagnostic of parallel substitution possibilities is discussed at length. All that the use of transformations as probes for constituent structure represents is a systematization and extension of the range of substitution classes which were understood to be relevant for identifying the membership of some word sequence as an instance of a particular category. It is, again, an argument based on common distributional properties.

7. This essential reference to the history of a given substring is not immediately apparent in *Syntactic Structures*, where the statement of the structural analysis uses a flat, not hierarchical, representation; but the latter is inherent in the way the structural analysis identifies the target word strings for the application of each rule. Phrase markers, to which transformational rules make reference, thus allow substrings of the sentence to be identified as having a single ancestor that dominates them exhaustively, thereby defining the movable, replaceable, or deletable material in the “structural change” part of the transformation.

We thus have a somewhat curious situation. On the one hand, the transitivity of domination inherent in constituency created the “structural history” that was essential for the application of transformational rules, leading Chomsky to adopt the Structuralists’s phrase marker notation (e.g., (2)) that permitted a transparent expression of those domination relations. The result was a data structure which made it possible to trace the sequence of domination relations holding for any terminal sequence arbitrarily far back into the rule expansion history licensing that sequence. On the other hand, transformational evidence was increasingly viewed as an essential part in validating specific phrase structure representation, based on the premise that if a string of words showed up in structural positions identified as solely the product of transformations (e.g., subjects of passive verbs), they were indeed necessarily to be regarded as syntactic constituents. The relationship between constituency and transformation thus verged on, and possibly into, circularity. The applicability of transformations, in itself, only defines objects to which the transformations are respectively applicable. Without some independent motivation for the increasingly intricate internal geometry attributed to the internal configurational structure of sentences via the rules of the base, the yield of these rules appeared to be doing little work, other than codifying the distribution of substrings labeled NP, AP, and so on. justified, in large part, on the basis of these same rules.

Perhaps not surprisingly, most of the explicit arguments on behalf of hierarchical constituent structures for syntactic representations are found in linguistic textbooks. One of the most extensive such defenses is given in Radford (1981), who provides the following distributional generalizations as the evidence that reference to such representations are necessary to account for the grammatical behavior of a word string:

1. recurrence of that string over a range of environments
2. coordinability
3. resistance to interposition of parenthetical material
4. proform replaceability
5. omissibility, for example, in ellipsis constructions

However, these criteria are plainly inadequate as either necessary or sufficient diagnostics for canonical phrase structure constituency. So far as sufficiency is concerned, we examine a wide range of cases in the following chapters in which both coordinability and omissibility fail as criteria for assigning the status of constituent to a word string.⁸

8. Parenthetical intrusion certainly fails as a criterion in English, as illustrated by the examples in (i) (where (ia) is a case of parenthetical intrusion into NP, and (ib), into PP):

- (i) a. I can’t find the, whatdya call it, SOCKET wrench that I need for this repair job.
- b. John contributed money to, if you can believe it, the Pythagorean Society!!

Nor, as Dalrymple (2001) notes, are all of Radford's criteria necessary to establish PS constituency. For example, we can hardly require that candidates for such constituent status be proform-replaceable, since there are many examples of accepted constituents which have no proforms (e.g., the canonical PPs *to John, of Mary*).⁹ In the absence of much more sharply defined and precise versions of these criteria, they do nothing more than, at best, give the analyst some rules of thumb—in effect, hunches—for carrying out preliminary probes into novel data. But they can hardly be accorded the status of “just in case” criteria as Radford seeks to do. On the whole, then, most of Radford's criteria do not actually establish anything about constituency, because they do not yield the intended results.

The substitution test that Radford proposes is somewhat different, as are most of the diagnostics suggested by Dalrymple (2001).¹⁰ This type of argument is highly effective as a defense of the claim that when word sequences combine with each other, the resulting longer word sequence is of a certain syntactic category, which can be referred to in stating the further combinatoric possibilities of that sequence; that is, one might argue on the basis of them that syntactic constituents corresponding to strings of words exist, with properties that are embodied in the category labels assigned to them, and with particular distributional properties.

This is actually a position that no major theoretical frameworks dispute. Strings of words may indeed instantiate some abstract syntactic type, such as NP or S, as opposed to being nothing more than lexical beads on a string, so to speak. But affirming the existence of such descriptions for strings of words in no way commits one to the position that when a pair of category-labeled strings *A, B* is combined by the composition rules of the syntax, the labels are retained in a hierarchically organized representation of a syntactic object that rules and constraints of the grammar can refer to. It is this latter, much stronger claim which is explicit in the representation of sentences as phrase markers, and so far as we can tell, none of Dalrymple's examples, or similar ones that have been offered as evidence for constituent structure, have any direct bearing on this

9. Conversely, the replacement test is not sufficient either: the \bar{N} proform *one(s)* and the VP proform *do so* are known to pick out discontinuous strings as antecedents, to display significant mismatches with their supposed antecedents, to appear without any actual antecedent in discourse, and so on; see, for example, Miller (1991), Ward and Kehler (2005), Houser (2010), Miller and Pullum (2013). We return to the issue of identity conditions in ellipsis in chapter 6.

10. Dalrymple mentions a number of phenomena, such as the distribution of the English possessive clitic, the possible occurrences of extracted English *wh* phrases on the left periphery of clauses, and the linear position of verbal adjuncts in Icelandic, which can be captured only by making reference to the syntactic category types of the word sequences with which the linguistic expressions in question combine.

claim.¹¹ In other words, a defense of constituency does not in itself constitute a defense of phrase structure hierarchy, notwithstanding the fact that many defenses of the latter in the history of syntactic theory have amounted to little more than arguments on behalf of the former. What was (and is) missing from such arguments is a demonstration that one actually needs the extra information about the history of the branching from higher nodes in phrase marker representation. Facts about selectional possibilities, in and of themselves, do not supply such a demonstration.

For this reason, the inventory of restrictions on *wh* movement presented in Ross's (1967) MIT dissertation appeared to be a watershed in the development of the still-new framework.¹² What Ross's analysis appeared to establish was the fact that certain phrase structural configurations acted in effect as trip wires blocking the normal movement of *wh* phrases to their ordinary positions on the left periphery of root clauses. Consider, for example, the contrast in (6):

- (6) a. Our research led us to believe that John was right to claim that matter could exist in ultra-condensed states.
 b. Which states did your research lead you to believe that John was right to claim that matter could occur in __?
 c. *Which states did your research lead you to believe John's claim that matter could occur in __?
 d. *Which states did your research lead you to the belief that John was right to claim that matter could occur in __?

At the time Ross's dissertation appeared, there was almost unanimous opinion in the field that the contrast between the judgments of well-formedness for (6b) on the one hand and (6c,d) on the other reflected facts about syntactic form. What Ross's account of such phenomena seemed to establish was that the syntactic failure responsible for the badness of the latter was a configurational condition, the Complex NP Constraint (CNPC), definable as a prohibition on the intervention of the configuration [NP . . . [S . . . , at any point on the extraction pathway between the *wh* extraction site and its surface position. In each of the ill-formed examples in question, there is a violation of the CNPC:

11. See the discussion of Generalized Phrase Structure Grammar (GPSG) in the next section in connection to this. Though GPSG took the notion of phrase structural constituency as part of its basic underlying architecture, due to its insistence of strict locality, the theory never made use of direct reference to hierarchical representations of sentence structures in stating the licensing conditions for well-formed sentences.

12. The rate of genuine progress in the early phase of transformational grammar was impressively rapid, and for those who went through the period it still seems a bit remarkable that Ross completed his thesis only a decade after the appearance of *Syntactic Structures*.

- (7) a. *Which states did your research lead you to believe [NP John's claim that [S matter could occur in __]]?
 b. *Which states did your research lead you to [NP the belief that [S John was right to claim that matter could occur in __]]?

In (7a), the critical boundaries occur closer to the extraction site than in (7b), but this doesn't matter: no matter where on the pathway this configuration occurs, moving the *wh* phrase past it results in the perception of major defectiveness. It follows that the CNPC, along with all of Ross's other conditions on extraction, jointly constituted support at the most fundamental level for the need to posit a hierarchy-based architecture for syntactic representations. In particular, without the hierarchical information recorded in phrase markers, there appeared to be no way to define the point in the application of movement transformations at which a fatal extraction occurred. During most of the history of transformational grammar, the existence of these constraints, reformulated though they were over the decades following Ross's original work, was taken as a "smoking gun" argument for a completely hierarchical representation of syntactic constituency that alternative approaches simply could not rebut.

During the past two decades, however, the picture of island effects as entirely (or even largely) syntactic in origin has undergone a very deep and extensive rethinking. Chapter 10 discusses a more recent, increasingly influential view of Ross's and similar restrictions on extraction which derives not solely from syntactic form but from the interaction of syntactic form with various functional factors (processing complexity, discourse coherence, and so on). A reassessment of the sources of islandhood along these lines seems very likely to drastically reduce the degree of confirmation island effects supply for the position which takes hierarchical configuration as the correct data structure for syntactic form. In this new picture, the compositional history of a sentence may still have a role to play in the conditions it must satisfy to ensure well-formedness, but the nature of this role need not constitute motivation for an architecture which incorporates a full representation of that history. We revisit this issue in greater detail in chapter 10.

1.3 Phrase Structure Grammar without Transformations

One of the most interesting (and perhaps perplexing) developments in the middle phase of syntactic theory post-*Syntactic Structures* was the fragmentation of generative grammar into a number of rival approaches, most of which rejected appeal to syntactic transformations. In the 1970s, the first work in Relational Grammar began to appear, followed by Generalized Phrase Structure Grammar (GPSG) and Lexical Functional Grammar (LFG), and subsequently the emergence of Head-Driven Phrase Structure Grammar (HPSG). The work of Michael Brame (see in particular Brame 1976, 1978),

though it never led to the development of a distinct research community, was nonetheless highly influential as well (in particular, the critique of the transformational rule of Equi-NP deletion in Brame [1976]). What is striking, in retrospect, is the fact that apart from Relational Grammar, the major alternatives to transformational grammar continued to take for granted the use of phrase structure constituency.¹³ This is particularly noteworthy in the case of GPSG, which merits detailed examination as an illustration of the persistence of the hierarchical phrase structure model in the nontransformationalist era—even where, as we argue directly, none of the transformationalist assumptions motivating that model held.¹⁴

The grip that the IC model (inherited from the Structuralists via its centrality in the language of syntactic representation enshrined in the dominant transformationalist paradigm) held on even non-derivational approaches is well illustrated in the case of GPSG. Of all the non-categorial models of grammar which emerged after the hegemony of transformational grammar began to break up in the early 1980s, GPSG, as we argue in detail in what follows, made the fewest assumptions that depended on hierarchical representations of syntactic form. The fact that such representations were assumed from the outset in GPSG thus testifies to the degree to which such representation had become an unquestioned default in linguistic theorizing.

Essentially, GPSG consists of a set of rules of immediate dominance, stating what the daughters of any given category may be, but not their linear order, and a set of constraints on the projection of two-generation phrase markers—“local trees,” in the theory’s parlance—from these rules. The major impact made by GPSG on a theoretical landscape in which Chomsky’s Revised Extended Standard Theory constituted the default assumption of most syntacticians was as an existence proof that one could capture grammatical dependencies of both the local “relational” sort (e.g., passive and “raising” patterns) and unbounded dependencies (in particular filler-gap linkages) without transformations. In retrospect, however, one could argue that its most profound break with prior work was its adherence to the ethic of strict locality at its very foundations, in that it precluded at the threshold all global structural conditions constraining dependencies of the sorts just mentioned or other morpho-syntactic linkages such as agreement and case-marking. In particular, GPSG’s architecture made it impossible to refer directly to aspects of the representation of a clausal structure beyond the syntactic properties of

13. The appearance of Ades and Steedman (1982), representing a renewal of a categorial grammar tradition largely dormant after Lambek (1958), was certainly important, but for some time after its appearance the sociological impact of this line of work was far less marked than was the case for GPSG and LFG.

14. Similar remarks apply to LFG, though for somewhat different reasons. In taking both configuration and grammatical relations to be primitive in its architecture, LFG appeared to be a kind of hybrid of phrase structure grammar and Relational Grammar.

a mother category and its daughters. In effect, a grammar can refer to no other aspect of a syntactic representation than what can be encoded in a single phrase structure rule.

GPSG's radical locality had obvious consequences for the kind of global restrictions that were routinely posited in transformational grammar following Ross (1967). Ross's CNPC as he stated it, for example, had no literal interpretation in the theory of GPSG in Gazdar et al. (1985); it had to be rephrased as a completely local constraint on the possible feature arrays attached to the nodes in which the dedicated gap-licensing feature SLASH appeared. This constraint, interacting with a network of similarly local restrictions, had the same effect as Ross's CNPC. The strict, absolute locality that was in a sense the central axiom of GPSG's architecture thus ensured that no reference to any configurational geometry beyond a mother node and its daughter nodes was either necessary or possible. Likewise, the theory of coordination and of feature sharing in Gazdar et al. (1985) ensured that a SLASH specification on a mother had to be shared with all its daughters, with lexical heads systematically excluded from this feature sharing, thereby capturing the Coordinate Structure Constraint. In retrospect, the lesson implicit in GPSG's account of islandhood is that the importance of syntactic hierarchy was much overrated.

One consequence of this strictly local formulation of island conditions is that, in an important sense, it undermines the role of islands in motivating hierarchical structure in syntactic representation. By treating unbounded dependencies via purely local feature percolation within two-generation trees, and stipulating that a nominal head daughter and a clause containing a nonempty SLASH specification could not compose, the system in Gazdar et al. (1985) made it unnecessary to assume the kind of global hierarchical structure that supposedly syntactic restrictions such as Ross's constraints seemed to vindicate. If the payoff of phrase structure hierarchy, via the dominance relation, was to make it possible to carry out operations on strings of words, then GPSG, which permitted no such operations, had no need for it, and this meant that the supposed rebuttal to such a view implicit in Ross (1967) no longer held.

From this point of view, one might read Gazdar et al. (1985) as an existence proof that reference to phrase markers was unnecessary in licensing natural language sentences. This situation is thus almost exactly the same as that of the steps in a standard logical proof, where earlier steps that led to the formulæ ϕ_1, \dots, ϕ_n , which jointly allow one to infer ϕ_{n+1} play no role whatever in that inference step. In other words, the crucial aspect of phrase structure-based theories of syntax which distinguishes such theories from logical deductive calculi (including those in type logics, such as we introduce in chapter 2)—the freedom to incorporate conditions which apply over the entire history of a sentence's licensing—is completely proscribed in GPSG (recall the discussion of [alleged] evidence for phrase structural constituency in the previous subsection in this connection). It is thus rather easy to picture reformulating GPSG in a way very similar

to that of a certain family of early categorial grammars (the so-called Ajdukiewicz/Bar-Hillel [AB]) calculus). Yet despite this fact, GPSG maintained a formal interpretation of its rules in which the latter license phrase markers rather than, for example, logical inferences. Thus, while the system presented in Gazdar et al. (1985) breaks profoundly with almost every aspect of the research tradition that emerged from Chomsky's work in the 1940s and 1950s, it preserves at its core Chomsky's reliance on the ultimately American Structuralist conception of phrase structure.¹⁵ There would have been a very specific advantage in recasting GPSG as a logic of category types. The information content of the mother category in a local tree corresponds, broadly speaking, to the interaction between a head category and its nonhead sister, and indeed this was the point of the "type-driven translation" mapping from syntax to semantics developed in Klein and Sag (1985) and the semantics chapters in Gazdar et al. (1985), which, for the first time, linked a fully explicit theory of feature-based phrase structure to the fully explicit model-theoretic semantics of Montague (1973). This was a remarkable achievement, but the difficulty of implementing this mapping will be evident to anyone who studies these sources. The essential idea of type-driven translation, as presented in Klein and Sag (1985), was quite simple and appealing: in any local tree, the daughter constituents will correspond to semantical objects whose respective type assignments determine only one way in which they can participate in function/argument application, and the resulting denotation is the semantics of the mother category. The mother category will, however, be a daughter category in the "next tree up" and contribute its semantics to the yield of *that* tree, and so on.

But the extreme awkwardness of the syntax-semantics interface in the later GPSG literature is a sign that the phrase structure architecture in which the framework was cast was fundamentally ill-suited to the simple syntax-semantics translation that Gazdar et al. (1985) were rather heroically striving for. Despite their efforts to define a maximally natural semantic translation mechanism to link form and meaning, the dif-

15. The situation in HPSG, GPSG's immediate descendant, is somewhat different. In the "middle" phase of HPSG, represented by Pollard and Sag (1994), local domination relations are encoded in categories via the feature DAUGHTERS (DTRS), which appears in all phrasal constituents in a representation (apart from those dominating a gap site). It is thus possible to define a feature name/value pathway between a phrasal constituent, its daughters, its daughter's daughters, etc.—effectively making possible statements of global conditions (e.g., index identity) on what are, in effect, arbitrarily deep phrasal configurations. For example, Condition C of the HPSG binding theory is stated in terms of a relation, generalized o-command, which makes use of this possibility. Later versions of HPSG, however, modify the feature geometry to make this possibility in principle unavailable, making the formalism closer to that of GPSG.

Similarly, in LFG, use of the Kleene star operator in the specification of side conditions on phrase structure rules allows path identities to be stated over unboundedly large syntactic depths. Thus, one can, via the device of "functional uncertainty," identify an extracted *wh* element with the grammatical role assigned to some constituent in a deeply embedded clause, and this property of the architecture of LFG is also essential for the interpretation of generalized quantifiers (see, e.g., Dalrymple 2001, 139–143, 250 ff.).

ference between syntactic composition (via hierarchical constituency) and semantic composition (by functional applications)—a difference inherent in the nature of phrase structure syntax—created an irreducible divergence between the two lacking any analytically natural remedy. And this conceptual disconnection between the construction of form on the one hand and meaning on the other led directly to major analytic difficulties, perhaps most obviously the framework's failure to sponsor any well-defined characterization of what we might call nonconstituent coordination, discussed in detail in the present monograph. Thus, even though GPSG broke with the concept that phrase markers were the data structure of choice for syntactic representation, the persistent reliance on the characterization of constituency by reference to phrase structure rules led to a major impasse in its empirical coverage which was, so far as we are aware, never satisfactorily overcome.

1.4 Conclusions

It is tempting, then, to say that, in the break between GPSG (and LFG) and the mainstream transformational approach, one can see an ironic echo of the same continuity with the past that we have observed in connection with the development of transformational grammar as an outgrowth—and in some sense, the culmination—of American Structuralism. On the one hand, the radical locality of GPSG represents a notable break with a global conception of syntactic representation in transformational grammar, where transformational applicability, island constraints, and various kinds of command relations deemed essential for statements of anaphora possibilities and the like made reference to arbitrary syntactic depths essential. On the other hand, although phrase markers did no actual work in terms of the theory, GPSG still relied on a variant of context-free phrase structure rules as the combinatory machinery of the framework and the syntactic platform for the semantic interpretation of sentences. The notion of phrase structural constituency was carried over, essentially as an unquestioned given, to nontransformational approaches as the basis for syntactic representations (recall once again the quote from Dowty [1996b] that we started this chapter with), just as the notion of immediate constituency from the Structuralists's analytic practice was carried over into the theoretical foundations of transformational grammar. And here again, the supposedly fundamental break between derivational syntax and its monostratal offshoots concealed what, from a certain perspective, can reasonably be regarded as their deeper commonality.

But a still deeper break with the Structuralist analytic tradition had been available in the theoretical marketplace from the very beginnings of the transformational era. The approach to grammatical composition offered in Lambek (1958) made no use of phrase structure rules at all. Nor did it derive representations of clausal structure as the progressive construction of an increasingly complex branching structure which termi-

nates in a string of words, substrings of which correspond to descendants of the root node and an arbitrarily large number of intermediate nodes. Instead, Lambek's system licensed sentences in precisely the way a logical proof licenses a conclusion based on one or more premises. In place of phrase structure rules, what has become known as the Lambek calculus provided an authentic logic—a set of connectives and rules of inference governing their use, which allowed one to deduce sentences as, quite literally, *theorems* of the lexicon under the proof theory provided by the calculus. It turns out that a principled extension of Lambek's proof theory leads directly to a drastic simplification in the syntax-semantics interface, whereby seemingly partial constituents in fact are derivable as full constituents, with a well-defined interpretation that enters into semantic composition in exactly the same way as phrase structure grammar's canonical phrasal types. In the following chapters we outline this extended proof theory and show how it accounts, in the most straightforward way, for a wide range of coordination and ellipsis phenomena that pose very difficult challenges for phrase structure-based approaches.

This is a section of [doi:10.7551/mitpress/11866.001.0001](https://doi.org/10.7551/mitpress/11866.001.0001)

Type-Logical Syntax

By: Yusuke Kubota, Robert D. Levine

Citation:

Type-Logical Syntax

By: Yusuke Kubota, Robert D. Levine

DOI: 10.7551/mitpress/11866.001.0001

ISBN (electronic): 9780262360807

Publisher: The MIT Press


Published: 2020

The open access edition of this book was made possible by generous funding and support from Arcadia – a charitable fund of Lisbet Rausing and Peter Baldwin



The MIT Press

© 2020 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license. Subject to such license, all rights are reserved. 

The open access edition of this book was made possible by generous funding from Arcadia—a charitable fund of Lisbet Rausing and Peter Baldwin.



This book was set in Syntax and Times Roman by the authors.

Library of Congress Cataloging-in-Publication Data

Names: Kubota, Yusuke, author. | Levine, Robert, 1947- author.

Title: Type-logical syntax / Yusuke Kubota, Robert D. Levine.

Description: Cambridge, Massachusetts : The MIT Press, [2016] | Includes bibliographical references and index.

Identifiers: LCCN 2020000483 | ISBN 9780262539746 (paperback)

Subjects: LCSH: Categorical grammar. | Grammar, Comparative and general—Syntax.

Classification: LCC P161 .K83 2016 | DDC 415—dc23

LC record available at <https://lccn.loc.gov/2020000483>

ISBN: 978-0-262-53974-6