Derivational Order in Syntax: Evidence and Architectural Consequences[*]

COLIN PHILLIPS & SHEVAUN LEWIS
University of Maryland
{colin, shevaun}@umd.edu

Most formal syntactic theories propose either that syntactic representations are the product of a derivation that assembles words and phrases in a bottom-to-top and typically right-to-left fashion, or that they are not constructed in an ordered fashion. Both of these views contrast with the (roughly) left-to-right order of structure assembly in language use, and with some recent claims that syntactic derivations and real-time structure-building are essentially the same. In this article we discuss the mentalistic commitments of standard syntactic theories, distinguishing *literalist*, *formalist*, and *extensionalist* views of syntactic derivations. We argue that existing evidence favors the view that human grammatical representations are the product of an *implementation dependent* system, i.e., syntactic representations are assembled in a consistent order, as claimed by grammatical models that are closely aligned with real-time processes. We discuss the evidence for left-to-right syntactic derivations, and respond to critiques of a proposal that the conflicts between the results of constituency diagnostics can be explained in terms of timing.

1. Introduction

Standard generative grammars describe language in terms that appear distant from considerations of everyday, real-time language processes. To some this is a critical flaw, while to others this is a clear virtue. One type of generative grammar defines a well-formed sentence as a static, structured representation that simultaneously satisfies all relevant constraints of the language, with no regard to how the representation is assembled (e.g., Sag, Wasow, & Bender, 2003). Another type of generative grammar defines a well-formed sentence as a derivation, or sequence of representations, that describes how the sentence is gradually assembled, often including various transformations that move words or phrases from one position to another in a structure. In the most popular current version of the derivational approach, derivations proceed 'upwards', starting from the most deeply embedded terminal elements in the sentence, which are often towards the right of a sentence (e.g., Chomsky, 1995; Carnie, 2006). Such derivations tend to proceed in a right-to-left order, which is

---

probably the opposite of the order in which sentences are assembled in everyday tasks such as speaking and understanding. Since these theories make no claim to being accounts of such everyday processes, the discrepancy causes little concern among the theories' creators. Generative grammars are typically framed as theories of speakers' task-independent knowledge of their language, and these are understood to be distinct from theories of how specific communicative tasks might put that knowledge to use.

Set against this background are a number of more recent proposals that various linguistic phenomena can be better understood in terms of derivations that incrementally assemble structures in a (roughly) left-to-right order. One can evaluate these proposals based simply on how well they capture the acceptability judgments that they aim to explain, i.e., standard conditions of 'descriptive adequacy'. But it is hard to avoid the question of whether it is mere coincidence that left-to-right derivations track the order in which sentences are spoken and understood. It is also natural to ask how left-to-right derivations impact the psychological commitments of grammatical theories. Are they procedural descriptions of how speakers put together sentences in real time (either in comprehension or in production)? Do they amount to a retreat from linguists' traditional agnosticism about 'performance mechanisms'? These are questions about what a grammatical theory is a theory of, and they are the proverbial elephant in the room in discussions of left-to-right derivations in syntax, although the issues have not been explored in much detail. Here we summarize the current state of some of the evidence for left-to-right derivations in syntax, and how this relates to a number of findings by our group and others on the nature of real-time structure building mechanisms. Some of these questions have been aired in previous work (e.g., Phillips 1996, 2004), but we have come to believe that the slogan from that earlier work ("the parser is the grammar") is misleading in a number of respects, and we offer an updated position here.

We start, however, with some important preliminaries on the question of what it is that grammatical theories aim to describe.

2. (Psychological) Goals of Grammatical Theory

Most contemporary linguists consider their theories to be theories of mental phenomena (and those that do not are of no concern to us here). What is less clear is what specific mental phenomena grammatical theories aim to describe. A standard response to this question is that grammatical theories aim to describe what a competent native speaker *knows* when he knows a language. This is often extended to cover the knowledge that a child brings to the language learning task, ensuring that he can arrive at the adult state. This is all very well as far as it goes. In the adult state, the knowledge in question is whatever-it-is that underlies a speaker's ability to reliably classify sentences as acceptable or unacceptable. But this mission statement leaves us with little guidance on how to interpret the specific components of the theories that we encounter. When we are told, for example, that the wh-word *what* is initially merged with a verb and subsequently moved to a left peripheral position in the clause, what claim is this making about the human language system? When this is described as part of a 'computational system', does this mean that there is a mental system that explicitly follows this sequence of operations? In our experience this is not a question that grammarians are typically eager to discuss, but as far as we can tell the answers to the question fall into roughly three types.

First, bottom-to-top 'cyclic' derivations may be understood as literal descriptions of sequences of mental operations that speakers may, in principle, carry out. They are just unlikely to do so in the normal course of speaking and understanding. This **literalist** position amounts to the claim that the grammar is one among a series of structure building systems that competent speakers have, distinct from but in some fashion related to structure-building systems for comprehension and production. It is not easy to find clear endorsements of this position, although it is a perfectly clear

12

cognitive hypothesis. We think that this is roughly what Townsend and Bever (2001) have in mind for their dual system model of sentence parsing, which analyzes all sentences using two different syntactic systems.

A second possibility is that the sequence of steps in a grammatical derivation may be understood not as a temporal ordering of structure-building operations, but rather as a formal relation between a set of structural representations. Under this **formalist** view, the full-blown representation of a sentence is a set of structural representations that are formally related to one another by the fact that they could, in principle, be sequentially ordered as a derivation. Speakers may even construct this full-blown set of representations in the course of everyday speaking and understanding, e.g., as the output of the comprehension process. But this view does not entail that the derivation describes the temporal sequence of steps that are used to assemble the full-blown representation. This is another position that we have not found frequent endorsements of in print, but it is certainly one that we have encountered in our discussions.

The third possibility is that a grammar is merely an abstract characterization of a function whose extension is all and only the well-formed sentences of a given language. Under this **extensionalist** interpretation of grammatical theory, a bottom-to-top grammatical derivation is not a hypothesis about a sequence of representations that a speaker would ever mentally construct, on any time scale. Consequently, theories that adopt such derivations should not be evaluated based on their correspondence with the mental computations that speakers actually carry out. Instead, such theories are accountable only to how successfully they partition the grammatical and ungrammatical sentences of a language. We suspect that this is a position that many practicing syntacticians are comfortable with, and it certainly corresponds with much standard practice in linguistics.

The literalist and formalist positions are clear mentalistic hypotheses about grammatical derivations, and so they are amenable to empirical scrutiny based on what speakers actually do. We will have more to say below about what speakers do. On the other hand, the extensionalist position is more elusive. By claiming that it is merely an abstract characterization of a function that generates the grammatical sentences of a language, it places itself beyond the reach of most empirical evidence, aside from acceptability judgments. If one takes this position seriously, then the individual components of a grammatical theory should be understood as having no independent status as mental objects or processes, as they are merely components of an abstract function, rather than components of a more concrete description of a mental system. Despite this, we encounter many extensionalist theories nowadays that appeal to notions of "efficiency" and "computational economy" of derivations, in ways that are hard to reconcile with the notion of an abstract functional description. In fact, the extensionalist position, when taken literally, undermines even the simple notion that children learn the individual parts of their target grammar, since that would amount to treating the components of the functional description as if they are actual mental objects. Of course, one is always free to avoid these consequences by moving towards a literalist or formalist position, accepting the additional mentalistic commitments that those positions entail.

Of greatest relevance to our current concerns is the relation between the goals of extensionalist grammatical theories and the putative attractions of left-to-right structure building mechanisms. If the goal is simply to describe an abstract function that generates all of the grammatical sentences of a language, with no regard to how speakers might actually generate sentences in real time, then it may be considered irrelevant to compare the grammar's derivations to the operations of comprehension or production systems. One could even argue that the comparison is misleading, since it amounts to a comparison of theories at two different levels of description, like comparing mathematical theorems with a pocket calculator. Under this view, it would be inappropriate to seek grammatical theories whose derivations resemble actual

13

mental computations, since that would go beyond what a grammatical theory is supposed to do: grammatical theories with bottom-to-top or left-to-right derivations, or with no derivations at all, are accountable only to their success in partitioning the good and bad sentences of a language. In the well-known terms of Marr (1982), the extensionalist position is that a grammatical theory corresponds to a description of the language system at the computational level (as opposed to lower level algorithmic and implementational descriptions).

In practice, the relevance of real-time processes to extensionalist grammatical theories varies, because linguists have different motivations for seeking abstract functional descriptions. For what we will call a **strategic extensionalist** the description of an abstract system that generates all of the sentences of a language is not an end in itself, but it is a reasonable interim goal, and one that can be pursued relatively easily and efficiently for a wide range of languages. For the strategic extensionalist, the aim is to ultimately move beyond the functional description to a more detailed, mechanistic understanding of the human language system. Consequently, the strategic extensionalist has every reason to be interested in constructing a theory that lends itself to progressing towards this next step. On the other hand, for a **principled extensionalist** the extensionalist enterprise is an end in itself, which remains relevant even if others are able to provide lower-level characterizations of the human language system. For the principled extensionalist it is important to distinguish *what* the human language system computes – this is viewed as the linguist's task – from *how* speakers carry that out in practical situations, which is something for psycholinguists to figure out. For the principled extensionalist, it is interesting if a left-to-right derivation happens to yield better coverage of acceptability judgments, but any resemblance between such derivations and real-time processes is irrelevant. In our view, the appropriateness of the strategic or the principled extensionalist approach depends on the question of whether the human language system is implementation (in)dependent.[1]

3. Implementation (In)dependence

In describing a biological or artificial system it is often useful to describe what the system does in abstract terms, and this is precisely what generative grammars aim to do. When a digital computer carries out a simple arithmetical calculation, it is certainly useful to be able to describe the calculation in abstract terms (e.g., $48 \div 8 = 6$) without reference to the algorithm that the computer uses to solve division problems or to the electronics of the chip that implements that algorithm. Similarly, it is useful to describe the contrast between English and Chinese wh-question formation in abstract terms (e.g., "English wh-phrases appear at the left edge of the interrogative clause, whereas Chinese is a wh-in-situ language"), setting aside questions of how English and Chinese speakers actually go about constructing wh-questions in real time, and ignoring the details of the neural processes associated with question formation.

We certainly do not dispute the usefulness of abstract descriptions, but it is important to distinguish between abstractions that are implementation independent and abstractions that are implementation dependent. An abstraction is **implementation independent** if the exact same abstract system can be realized in different ways by multiple lower-level implementations, with no change in the abstract system itself. Mathematical systems are straightforward examples of this: simple arithmetical calculations may be implemented using a variety of equivalent algorithms and using a host of different types of hardware (digital computer, abacus, human brain, etc.), with no impact at all on the nature of the arithmetical facts. In contrast, an abstraction is

---

[1] Jackendoff (2002, ch. 2) offers a useful discussion of these issues with an added historical perspective. Jackendoff distinguishes between 'soft' and 'hard' idealizations.

14

implementation dependent if it is only ever realized in one way at a less abstract level of description.

Abstract generative grammars have traditionally been studied under the assumption that they are implementation independent. It is assumed, for example, that eventually we will be able to write software or build hardware to implement the human language system in computers that differ from human 'hardware' in many respects. As such, this is taken as evidence for implementation independence. This, however, is clearly an empirical question. It is a real and attractive possibility that computational modelers of the future might succeed in that effort, but it is far from having been accomplished at this time. The fact that some grammars that approximate fragments of human language have been digitally implemented does not settle the question. A relevant model should minimally be completely descriptively adequate, and one might reasonably want the model to approximate human linguistic abilities in more ways than simply classifying (un)acceptable sentences. Such models currently inhabit the realm of thought experiments. Here we disagree with Neeleman & van de Koot (2010), who agree with us on the importance of implementation (in)dependence, but take it as quite significant that some natural language grammars could be implemented in different artificial devices, concluding that human language is implementation independent.[2]

In the meantime, we argue that the emphasis on the implementation independence of generative grammars is misplaced if the purpose of the endeavor is to understand the nature of the human system. If a research community makes the choice to carve out one sub-area of the study of human language abilities, e.g., classifying (un)acceptable sentences, and then finds that this sub-area can be described in implementation independent terms, then this surely does not entail that the sub-area corresponds to a privileged, implementation independent sub-area of the human language faculty. There is no doubt that it is interesting and useful to develop explicit computational models of human language, but such models cannot show whether human grammatical abilities are, in fact, independent of their cognitive or neural implementation. Regardless of what one's personal hunch is about the likelihood of successful non-human implementation of language, a more interesting question for our purposes is whether it is implementation independent within humans.

According to the standard formulation of the principled extensionalist view, speakers of a language have knowledge of the (un)acceptable sentences of their language, and this knowledge can be described at an abstract level by a generative grammar. Importantly, it is assumed that this abstract knowledge is implementation independent, and that speakers can put their knowledge to use in different ways in activities such as speaking, understanding, and internally generated speech. This means that for any given well-formed sentence structure defined by the grammar a speaker may have a number of different ways of mentally assembling the structure. If this assumption is supported, then the principled extensionalist is justified in separating speakers' knowledge of what is well formed from their knowledge of how to assemble well-formed structures. But if speakers do not, in fact, have multiple ways of constructing

---

[2] Neeleman and van de Koot describe a thought experiment in which a pair of computers with radically different underlying architectures are both able to pass a kind of Turing Test for human language, such that their performance is indistinguishable from human native speakers of the language. They argue that one would want to say that both computers speak the relevant language, and conclude that this is because the computers and the humans would share an abstract, implementation-independent grammar. We agree with the intuition behind this thought experiment, but it seems to reveal more about the common usage of predicates like "speaks French" than about the nature of the human language faculty. Despite this misgiving, we find much to like about Neeleman and van de Koot's discussion of the challenges involved in understanding grammatical theories at different levels of description.

15

the same representations, then human language appears to be more implementation dependent.

The evidence on whether speakers have multiple ways of constructing the same sentence representation is not extensive, but there are a number of reasons to think that speakers have one and only one way of assembling any individual representation. There is almost no evidence for the alternative view that speakers have multiple ways of building the same representation. First, speaking and understanding proceed in the same (roughly) left-to-right fashion. Although they have different goals, they have a great deal in common, appear to construct the same representations, and plausibly do so in the same order, although this topic has not been investigated in great detail. Second, in comprehension and production there is much evidence that speakers build structures and interpretations incrementally, in roughly the order in which words are presented. We are unaware of evidence that speakers are able to construct the same interpretation in different orders. For example, reading backwards is a task that lies somewhere in the difficult-to-impossible range, despite its correspondence with the derivational order assumed in many generative grammars. Third, in cases of reanalysis in sentence understanding, where comprehenders are led into a 'garden path' and must then reorganize their initial parse, some evidence suggests that speakers repair by returning to the error point and simply re-parsing the sentence in exactly the same order that it was presented (Inoue & Fodor, 1995; Grodner et al., 2003).

We therefore adopt the working hypothesis that natural language grammars are implementation dependent with respect to how sentences are assembled: there is only one algorithm for structure-building across all human speakers (of the same language). We would certainly welcome more systematic evidence, but currently the evidence for the alternative implementation independent position is practically non-existent. The implementation dependent position is the simpler and more falsifiable hypothesis, and hence should be preferred until proven otherwise. Consequently, we think that the principled extensionalist position is unwarranted, and that the motivations for developing abstract generative grammars are more pragmatic than principled. We have no problems with pragmatic motivations, and we recognize the value of focusing on characterizing good/bad sentences as a way of making headway in describing human language. But this is very different from the position that the characterization of good/bad sentences is a fundamentally separate enterprise from understanding real-time language processes. Ultimately we seek theories that capture how sentences are put together, and not just what their final form is.[3]

In discussions of these issues we sometimes encounter an objection in the form of a 'slippery slope' argument. If human language is better described in terms of real time cognitive processes rather than abstract functions, so the argument runs, then why stop there – why not continue all the way down to the level of neurochemical processes in brain cells or beyond? As with other slippery slope arguments, the expectation is that we should find this consequence appalling, and hence should drop the entire argument. We acknowledge the concern, but disagree with this argument. First, we are not arguing against the usefulness of abstract descriptions in the study of language. We find them exceedingly useful. We are simply arguing that there is no privileged level of abstraction – the level occupied by most current generative grammars – that is

---

[3] Although we have focused our attention here on the interpretation of derivational grammatical theories, the issue of implementation (in)dependence is equally relevant to interpreting grammatical theories that assume no derivations. Non-derivational syntactic theories, which are often presented as preferable due to their order-neutrality, imply implementation independence. If syntactic representations are not, in fact, implementation independent, then that should count against non-derivational theories, just as it counts against bottom-to-top derivations.

exempt from considerations of psychological implementation. Second, it should be emphasized that implementation (in)dependence is not an all-or-nothing property of an abstract system. Rather, a system's implementation (in)dependence must be evaluated at each successive degree of abstraction. Our focus in the current discussion is on the relation between symbolic descriptions of the structure of sentences and symbolic descriptions of the procedures for assembling sentences in real time. This leaves open the possibility that human language is implementation independent at a lower level. For example, there are interesting arguments that the basic notion of hierarchical structure in language must be implementation independent with respect to its neural encoding, because of the very different demands of immediate vs. long-term encoding of sentences in memory (Jackendoff, 2002; van der Velde & de Kamps, 2006). Briefly, the most plausible method for long-term information storage in the brain – through changes in synaptic connectivity – is too slow for the millisecond-scale processes needed for real-time language use. The need for dual encodings of the same structures suggests that structured representations are not always neurally encoded in the same fashion, and hence that they are implementation independent. It is far from certain that this argument goes through, because it is not clear that immediate and long-term encodings of sentences in memory are genuinely isomorphic to one another. But this is the type of argument that might one day lead to a clear finding of implementation independence for human language (for one specific level of implementational detail). In sum, the slippery slope argument is misplaced in the current case.

We should also emphasize that we do not claim that a grammatical theory is inherently preferable if it can be transparently linked to real time processes. Standard considerations of descriptive adequacy (i.e., generating all and only the acceptable sentences of a language) are as relevant as ever, and all theories should be accountable to that standard. A grammatical theory that achieves impressive descriptive adequacy using mechanisms that are opaque to real time implementation is an interesting theory nonetheless. However, the descriptive success of such a theory begs for further analysis of why it works so well, and whether its success crucially depends on the opaque mechanisms. This analysis should ideally lead to development of a more psychologically transparent version of the theory. We take up this challenge below for arguments that have been presented in the syntax literature in favor of bottom-to-top derivations.

4. Psychological Aims of a Real-time Grammar

Having examined the psychological commitments of standard generative grammars, it is appropriate to apply similar scrutiny to grammars that adopt roughly left-to-right derivations and that aspire to be models of real time processes. Here it is important to address some possible misconceptions. (One of us bears some blame for spreading some of the misconceptions.)

First, the slogan 'the parser is the grammar' (Phillips, 1996) sounds nice enough, but it is unfortunately misleading, as it too closely identifies the grammar with the task of language comprehension. It would probably be more appropriate to regard a real-time grammar (the 'structure builder') as one important component of a parsing/comprehension system, but certainly not the whole system. We envision a structure building system that combines words and phrases to form sentences and meanings in real-time, in essentially the same manner in comprehension and production. This system is task-neutral, and it could even operate in the absence of external input or a communicative goal, i.e., neither comprehension nor production is necessary. The system is, of course, put to use in speaking and understanding, where the structure that it builds is constrained by the external input (comprehension) or by the speaker's message (production), but these tasks require far more than just a structure building system. Parsing requires some way of using incoming sounds and words to determine which specific structure building operations are appropriate.

17

Production requires some way of mapping the elements of an intended message to operations in the structure building system. It appears that the human parser is fairly good at mapping external input onto appropriate structure building operations, and the human production system appears to be similarly effective at selecting structure building operations that are appropriate to the speaker's communicative goals. But the assumption that natural language grammar has the form of a real-time structure building system does not logically entail the success of either parsing or production. Consequently, some caution is needed when using evidence from parsing or production to assess the nature of the structure builder.

Second, and closely related to the first point, the claim that the mental grammar should be understood as a real-time structure generation device does not guarantee that it is part of a perfect parsing device. When we claim that the real-time structure-building device is the mental grammar, we predict that the representations that this device constructs should be grammatically well formed, and that it should incorporate whatever detailed grammatical machinery we would normally expect of a grammar. This means that the system should not construct the rough-and-ready representations that some have argued to be created in the service of rapid, efficient comprehension (e.g., Ferreira & Patson, 2007; Townsend & Bever, 2001). However, the claim that the representations that are built on-line are grammatically precise does not entail that they are the same ones intended by the speaker. In extreme cases where a listener is distracted or in a noisy environment, he might use his mental grammar to construct a perfectly reasonable representation that bears only a weak relation to the sentence that the speaker uttered. This would show that the listener is failing to make full use of the input, but it would not license any clear conclusions about the nature of the listener's real-time structure building system. Consequently, when we examine studies of on-line satisfaction of grammatical constraints, the key prediction for current purposes is that *the comprehender constructs grammatically well-formed representations*, even if those representations are not a grammatically possible parse of the incoming sentence. We will have more to say on this point in a moment.

Third, the claim that the grammar has the form of a real-time structure building system is independent of long-standing psycholinguistic questions about how speakers resolve syntactic ambiguities in language comprehension. Structural ambiguities arise when the input to the comprehension system has two or more well formed structural analyses, i.e., they are cases where the grammar alone cannot decide. Ambiguity resolution has enjoyed a prominent position in psycholinguistic research on sentence comprehension, and there have been long and heated debates over which types of information are brought to bear in resolving ambiguity (e.g., Frazier, 1987; MacDonald, Pearlmutter, & Seidenberg, 1994; van Gompel & Pickering, 2007). There is, in fact, broad consensus that simpler parses are favored in ambiguity resolution. The controversy surrounds what it means to be 'simpler' (e.g., structurally simpler, more frequent, more semantically or pragmatically natural, etc.) and what it means to be 'favored' (i.e., the unique parse vs. highest ranked parse among multiple parses pursued in parallel). These discussions are interesting, but they are orthogonal to our claims about the form of the grammar.

Fourth, we have been surprised at how often claims about the procedural nature of the grammar are interpreted as claims that grammatical phenomena are epiphenomenal. In linguistics and psycholinguistics one often encounters proposals that some phenomenon that has traditionally been analyzed in formal terms can instead be explained in terms of constraints on parsing or production (e.g., Kluender & Kutas, 1993; Hawkins, 1994; Harris & Bates, 2003; Hofmeister & Sag, 2010). These reductionist accounts of linguistic phenomena are often referred to as 'processing accounts', as a way of contrasting them with 'formal' or 'grammatical' accounts. They are claims that the phenomenon in question does not fall under the purview of the mental grammar. Reductionist analyses of grammatical phenomena are interesting, and we find some types of evidence more compelling than others (for discussion see

18

Sprouse, Wagers, & Phillips, 2012; Phillips, 2013ab), but they are quite different than the models that we are advocating here. Since reductionist accounts are explicitly non-grammatical in nature, they clearly have little in common with the claim that grammatical derivations follow an order that happens to be well suited to comprehension and production.

Finally, we should clarify the reason why we repeatedly describe grammatical derivations as proceeding in a *roughly* left-to-right order. Sentences are spoken and heard in a strict left-to-right order. (Strictly speaking this is a tautology, as left-to-right order is merely a conventional representation of the temporal order of words in speech.) However, it is probably not the case that mental structure building operations perfectly follow the linear order of a sentence, whether in comprehension or production. To take just one example, in a head-final language such as Japanese it may be necessary for the structure building system to create a position for the head of a phrase before it has completed the arguments and adjuncts that precede the head. More generally, structure building in comprehension is probably not entirely synchronized with the appearance of words in the input. There is growing evidence that comprehenders often build structural positions in their parses before encountering the words in the input that phonologically realize those positions (Aoshima, Phillips, & Weinberg, 2004; de Long, Urbach, & Kutas, 2005; Lau et al., 2006; Mazuka & Itoh, 1995; for review see Lau, 2009), and some evidence for related effects in production (Momma, Slevc, & Phillips, 2013). The upshot of this is that it may not even be desirable to insist upon a strict left-to-right order for grammatical derivations, since the operations of the real-time structure builder may not proceed in a strict left-to-right order. If it is the case that there is a single structure-building system that assembles sentences in a strict order, then it is likely that this order will turn out to be only *roughly* left-to-right. What matters is whether the order of structure building operations is consistent.

5. Is On-Line Structure Building Grammatically Precise?

If the representations that are built during real-time comprehension and production differ from those that are motivated by standard grammatical analysis, then we face *prima facie* evidence for multiple structure building systems, and thus evidence against our claim of a single procedural grammar. Fortunately, there is now much experimental evidence that bears on this question, and it mostly supports the notion that real-time processes assemble syntactic representations that are the same as those motivated by grammatical analysis. We should reiterate here that the question of how grammatically sophisticated real-time structure building is must be kept distinct from the question of how accurately the parser analyzes the input. A parsing system could, in principle, build wonderfully rich grammatical representations that are nevertheless rather poor analyses of the incoming strings of words.

Many classic psycholinguistic studies from the 1960s and 1970s addressed the question of whether on-line language processes create representations with the detailed structural properties that were proposed in the newly emerging field of generative grammar. This body of work is most often remembered for its failure to find correlates of the transformational derivations found in theories of that time,[4] but a

---

[4] The actual findings from these studies were less definitive than is typically claimed in historical reports (e.g., Townsend & Bever, 2001, ch. 2). The tests focused on a rather narrow linking hypothesis, and many of the conclusions relied on specific grammatical analyses that have not stood the test of time. Tests of the 'Derivational Theory of Complexity' (DTC) examined the narrow linking hypothesis that the perceptual complexity of a sentence is best predicted by the number of steps in its transformational derivation. But given the many different factors that contribute to the perceptual complexity of a sentence, there is little reason to expect that derivation length should be the primary

more enduring outcome is that the same studies often found experimental support for the surface structure representations of the time (for reviews see Fillenbaum, 1971; Fodor, Bever, & Garrett, 1974; Levelt, 1974). For example, hierarchical clustering analyses of speakers' relatedness judgments for word triads from a sentence yielded a good approximation to the surface structure of a sentence (Levelt, 1970), and studies on the perceptual mis-location of non-speech clicks played during a sentence confirmed a difference between object control sentences like *The general defied the troops to fight* and exceptional case marking constructions like *The general desired the troops to fight* (Bever, Lackner, & Kirk, 1969).

Far more recently, studies that use highly time-sensitive measures such as event-related brain potentials (ERPs) have made it possible to track how quickly comprehenders are able to detect different types of anomaly in the linguistic input. This work has shown that speakers detect just about any linguistic anomaly within a few hundred milliseconds of the anomaly appearing in the input. Different types of grammatical anomalies elicit one or more from among a family of different ERP components, including an (early) left anterior negativity ('(e)LAN'; Neville et al., 1991; Friederici, Pfeifer, & Hahne, 1993) or a P600 (Osterhout & Holcomb, 1992; Hagoort, Brown, & Groothusen, 1993). Many questions remain about what the different components reflect and what determines which components are evoked in any individual situation (Hagoort, 2003; Friederici & Weissenborn, 2007; Lau, Phillips, & Poeppel, 2008; Federmeier & Laszlo, 2009; Gouvea et al., 2010; Brouwer, Fitz, & Hoeks, 2012), but for current purposes the most relevant outcome from this research is that more or less any grammatical anomaly elicits an ERP response within a few hundred milliseconds. If the on-line analyzer is able to immediately detect any grammatical anomaly that it encounters, then it is reasonable to assume that it is constructing representations that include sufficient grammatical detail to detect those anomalies.

Another body of on-line studies has examined whether on-line structure building respects various grammatical constraints, i.e., whether the parser ever creates grammatically illicit structures or interpretations. Many studies have found evidence of immediate on-line effects of grammatical constraints, such as locality constraints on wh-movement (Stowe, 1986; Traxler & Pickering, 1996; Wagers & Phillips, 2009), and structural constraints on forwards and backwards anaphora (Kazanina et al., 2007; Nicol & Swinney, 1989; Sturt, 2003; Xiang, Dillon, & Phillips, 2009; Lewis, Chow, & Phillips, 2012; Dillon et al., in press). These findings extend to complex cases that present apparent challenges for incremental application of grammatical constraints, such as constraints on backwards anaphora in Japanese, where the constraints must apply before any verb has appeared in the input (Aoshima, Yoshida, & Phillips, 2009), and constraints on parasitic gaps inside complex subjects in English, where the parasitic gap precedes its licensor (Phillips, 2006). Findings such as these imply that the structures created on-line include sufficient structural detail to allow the constraints to be applied during parsing.

The many different types of evidence for on-line grammatical sensitivity do not, of course, strictly require that the structures that are built on-line are exactly those that

---

predictor of processing difficulty. Nevertheless, early evidence based on some transformations that have stood the test of time, such as passivization and subject-auxiliary inversion, provided surprisingly good support for the DTC's predictions, and these findings were not effectively challenged, despite reports to the contrary (for review see Phillips, 1996, ch. 5). Subsequent failure to find perceptual complexity increases caused by such operations as the transformation that converts a full relative clause ("the house that was red") into a pronominal adjective ("the red house") are probably not surprising. We would not want to argue that the DTC was substantiated, but the reports of its defeat strike us as somewhat stylized history.

are sanctioned by the grammar, or that the system that builds them is a procedural grammar. It is always possible that the on-line structure builder is not, in fact, identical to the grammar, but instead is a very effective 'covering grammar' for the true grammar of the language. But to the extent that the on-line structure builder constructs exactly the right representations to capture both on-line behavior and standard acceptability judgments, then we see little motivation for postulating an independent grammar that yields no additional empirical coverage.

There are, however, a number of findings in the psycholinguistics literature that have been taken to indicate divergence between the structures created on-line and those motivated by traditional grammatical analysis. We will briefly review three types of evidence, from misinterpretations, syntactic priming, and grammatical illusions.

Comprehenders frequently misinterpret the sentences that they encounter. Fernanda Ferreira and colleagues have argued that this is a desirable property for a comprehension system, and that the misinterpretations are the result of 'good enough' (GE) representations, which contrast with representations that are "detailed, complete, and accurate with respect to the input" (Ferreira & Patson, 2007, p. 71; see also Ferreira, Bailey, & Ferraro, 2002). One recent event nicely illustrates the intuition behind Ferreira's argument. On October 2nd 2008 a record TV audience (70 million in the US alone) watched the debate between the vice presidential candidates, Senator Joseph Biden of Delaware and Governor Sarah Palin of Alaska. In a segment on energy and climate change Palin said, "I'm not one to attribute every activity of man to the changes in the climate". Probably most viewers immediately recognized what Palin intended to convey with this quote (that climate change is not primarily caused by human activity), and few would have even noticed that the literal interpretation of her statement is quite different (that human activities are not exclusively caused by climate change). Palin's slip-of-the-tongue was (roughly) a classic exchange error, and comprehenders successfully recovered Palin's message despite the error, presumably because they had a good idea of what the message was likely to be. Ferreira and colleagues argue that situations like this show GE representations in action, and straightforwardly illustrate one of the benefits of an interpretive system that is not a slave to the precise surface form of incoming sentences.

Ferreira and colleagues offer additional evidence of systematic misinterpretations from experimental studies. For example, they tested what interpretations speakers take away from garden path sentences in which they initially misparse and must then reanalyze, e.g., *While Anna dressed the baby played in the crib* (Christianson, Hollingworth, Halliwell, & Ferreira, 2001). When questioned after the sentence, speakers reliably agreed that the sentence stated that the baby had played, but they also agreed on a substantial proportion of trials that the baby was dressed, suggesting that in their final interpretation the single NP *the baby* filled two conflicting thematic roles. They refer to this as a 'lingering' garden path effect (see also Sturt, 2007). Similarly, Ferreira reports that comprehenders often incorrectly judge that the surface subject is the 'do-er' in passive sentences like *The dog was bitten by the man* (Ferreira, 2003).[5]

We agree with Ferreira that the misinterpretation of good sentences and the successful repair of speech errors are important phenomena that show that on-line interpretation does not always deliver the correct meaning of an incoming sentence. But as we have

---

[5] In this study comprehenders were probably not simply following a plausibility-based heuristic, as they also showed very similar error rates in judgments of non-reversible passives, such as *The cheese was eaten by the mouse*. Only in the most difficult sentence types (object clefts) did error rates increase in the manner predicted by a plausibility-based heuristic.

emphasized above, there is a difference between failure to correctly analyze the incoming string of words and failure to build a grammatically well-formed representation in response to the input. As psycholinguists we are certainly very interested in the question of how effectively people parse, but this is separate from the question of whether the on-line structure builder generates only grammatically well-formed representations. The evidence reviewed by Ferreira does not show that comprehenders assemble grammatically ill-formed representations. We suggest that in the case of 'lingering interpretations' of garden path sentences what happens is that comprehenders incrementally update their interpretations over the course of the sentence, but that interpretations are not labeled based on the pieces of syntax that generated them. Consequently, when a syntactic parse undergoes reanalysis to a complete and well-formed structure, rescinding of the incorrect syntactic analysis does not lead to rescinding of any interpretations that were previously generated by that syntactic analysis. Interpretive repair of speech errors, as in the example from the vice presidential debate, presents a slightly different situation, since the repairs occur in situations where the comprehender has a good idea of the intended meaning and hence the speech input is redundant. In such cases the interpretive system likely completes its task before the syntactic parse has finished. It is interesting that the interpretive system is able to do this,[6] but it tells us little about the nature of the syntactic representation that is generated in these situations, which may be perfectly grammatically well-formed. A related observation that may be relevant here is that different types of speech errors appear to impact language comprehension in different ways. In sentences that are syntactically appropriate but semantically garbled, comprehenders are sometimes able to recover the intended interpretation without even noticing the error, as in the example above. In contrast, when speech errors lead to syntactic deformation of a sentence, comprehenders are generally able to recover the intended form, but the error does not pass unnoticed. This suggests that comprehenders may be able to skip detailed interpretation of an incoming sentence, but that they cannot skip syntactic analysis, even if it is not needed for interpretation.

A second potential argument for a mismatch between on-line structure building and the structures sanctioned by the grammar comes from syntactic priming. In language production, many studies have shown that the use of a particular syntactic structure increases the likelihood that the same structure is used in subsequent utterances, even when there is no lexical overlap between the utterances that have the same structure (Bock, 1986; Pickering & Ferreira, 2008). This finding has led to interesting questions about what pairs of structures count as the 'same' for purposes of syntactic priming. An influential study by Bock and Loebell (1990) showed priming between VPs containing PPs that differ in their thematic roles, e.g., VPs with a locative PP such as *The wealthy widow drove the Mercedes to the church* primed VPs with a recipient PP such as *The wealthy widow gave the Mercedes to the church*, and VPs with a locative *by*-phrase such as *The 747 was landing by the control tower* even primed full passives such as *The 747 was alerted by the control tower*. (Note that we use lexically matched examples here for illustrative purposes only. The adjacent prime-target pairs in these studies did not overlap in this way.) In addition, many other studies have shown that syntactic priming does not depend on overlap in the function words (including prepositions) between primes and targets (Bock, 1989; Fox Tree & Meijer, 1999). Taken together, the evidence therefore suggests that relatively coarse-grained syntactic parallelism is sufficient to cause structural priming. This must somehow be

---

[6] This account clearly begs the question of how the relevant interpretations are accessed or generated, if not through detailed compositional interpretation of the sentence structure. Investigation of this question may well lead to the conclusion that the language comprehension system has multiple ways of using cues to generate interpretations. However, this is quite different from the conclusion that there are multiple representational systems or multiple real-time structure building procedures.

reconciled with the evidence that motivates linguists to postulate fine-grained structural distinctions between superficially similar sentences. Different responses to this challenge are possible, and they remain to be resolved. The structural priming evidence could be taken as support for grammatical models that make less fine-grained structural distinctions (Culicover & Jackendoff, 2005). Or it could be taken to show that the representations involved in structural priming are not those defined by the grammar (we do not favor this option, but it is a logical possibility). Alternatively, it could be taken to show just that the structural priming paradigm is a relatively blunt tool for investigating structure, because relatively coarse-grained similarity between structures is sufficient to cause structural priming. This would leave open the possibility that on-line processes build fine-grained structures.

A third potential motivation for distinguishing the structures built on-line and those sanctioned by the grammar comes from grammatical illusions, cases where comprehenders appear to fleetingly accept structures that are judged bad after more reflection. The most notorious case of a grammatical illusion involves comparative constructions such as *More people have been to Russia than I have*, which are semantic gibberish but initially sound remarkably good.[7] Townsend and Bever (2001) argue that such cases make an interesting case for a system that distinguishes a rough-and-ready initial analyzer from the more fine-grained analyses of the grammar. A number of additional cases of illusory acceptability have emerged in recent years. One case involves the spurious licensing of the negative polarity item (NPI) *ever* by non c-commanding negation, as in *The bills that no senators have supported will ever become law* (Drenhaus, Saddy, & Frisch, 2005; Xiang, Dillon, & Phillips, 2009). Another case involves illusory agreement licensing in which a plural-marked verb is judged to be acceptable in the vicinity of a plural NP that is not its syntactic subject, as in *The runners who the driver see ...* or *The key to the cabinets probably are ...* (Clifton, Frazier, & Deevy, 1999; Pearlmutter, Garnsey, & Bock, 1999; Wagers, Lau, & Phillips, 2009). Yet another case involves evidence that during the processing of pronouns comprehenders fleetingly consider a clause-mate subject NP as a potential antecedent, in violation of Binding Principle B (Badecker & Straub, 2002; Kennison, 2003), although findings are mixed, and a number of other on-line studies report immediate effects of Principle B (Clifton, Kennison, & Albrecht, 1997; Lee & Williams, 2006; Nicol & Swinney, 1989; Runner, Sussman, & Tanenhaus, 2006; Lewis et al., 2012). For a more detailed review of where grammatical illusions do and do not arise see Phillips, Wagers, & Lau (2011).

Grammatical illusions pose a challenge to our hypothesis of a real-time procedural grammar to the extent that they justify a mismatch between on-line structure building processes and processes that operate under less time pressure. However, demonstrations of illusory acceptability do not automatically show such a mismatch. Grammatical illusions could have a number of sources other than the existence of multiple structure building systems. They could reflect mis-parsing, where the comprehender builds a perfectly grammatical representation that happens to not match the input. Alternatively, they could show that the normal workings of the structure building system are slow enough that on-line methods can probe the intermediate steps of the computation. Wellwood and colleagues have recently argued that illusory comparatives reflect mis-parsing rather than construction of an ill-formed internal representation. In acceptability rating studies they find that the illusions are more robust with predicates that are 'repeatable' such as *go to the gym* than with predicates that are 'non-repeatable' such as *won the lottery yesterday*, and suggest that this is

---

[7] These comparative illusions, which some have referred to as *Escher sentences*, were first pointed out by Montalbetti (1984), although he presents them as a curiosity and does not offer an analysis of the phenomenon, which has not been systematically studied until recently.

because the sentences are mis-interpreted as event quantification rather than as quantification over individuals (Wellwood et al., 2013). Notably, English grammar allows the syntax of individual quantification to be interpreted as event quantification in certain contexts, as in *The Washington DC metro carries more than 200 million passengers per year*, which is a claim about person-trips rather than about distinct individuals.

The other cases of grammatical illusions that we have listed here are probably not amenable to a mis-parsing analysis, but might nevertheless reflect the normal operations of a grammatically accurate structure building system. The temporary consideration of local antecedents for pronouns that violate Principle B (as found in some, but by no means all on-line studies of Principle B) may be directly related to a natural grammatical implementation of the constraint, applying as a filter that marks candidate referential dependencies as illicit, rather than as a constraint that prevents the generation of illicit candidates. Since pronouns may be associated with a wide range of syntactic and discourse antecedents, a mechanism that generates candidates and then excludes inappropriate candidates may be the most feasible way of applying Principle B, and it may be that the fleeting consideration of illicit antecedents found in some studies reflects this mechanism in action. More broadly, 'generate-and-filter' mechanisms are familiar from many grammatical theories (e.g., Chomsky, 1981; Legendre, Grimshaw, & Vikner, 2001), and hence are plausible components of a real-time grammar. The case of illusory NPI licensing remains poorly understood, as few NPIs and only a narrow range of syntactic and semantic contexts have been tested to date, so it is unclear how far illusory NPI licensing extends beyond the specific environments examined so far. However, Xiang and colleagues have attempted to show how the illusions could reflect the operations of the same pragmatic licensing mechanisms that are widely held to be responsible for normal NPI licensing (Xiang et al., 2009). According to this account, the illusions reflect inappropriate pragmatic inferences, rather than use of rough-and-ready licensing mechanisms. In the case of illusory agreement licensing Wagers and colleagues have argued that the illusions arise from the use of grammatically fully appropriate feature retrieval operations in an architecture with noisy memory representations (Wagers et al., 2009). We should emphasize that systematic investigations of grammatical illusions are relatively new and that none of these accounts are definitive. However, they show that it is possible to account for the phenomena using mostly standard structure-building and interpretation mechanisms, without recourse to independent on-line and off-line systems. One clear advantage of the single-system accounts is that they predict that only certain types of temporary illusions should be possible. Accounts that invoke a separate rough-and-ready structure building system are less constrained.

6. Bottom-to-top derivations

We have argued that human grammars are implementation dependent, at least with respect to how syntactic structures are assembled, and that it is desirable and plausible to view grammars as procedural systems that can be understood in terms of actual real-time mental computations. However, most contemporary syntactic theories seem incompatible with this view: they assume that grammatical derivations proceed in a strictly bottom-to-top order, progressively combining words and phrases starting from the most deeply embedded elements. For most languages this yields derivations that proceed in a mostly right-to-left order, i.e., the opposite of the order in which comprehension and production operate. We must therefore address the evidence for bottom-to-top derivations, to determine whether it undermines our approach. We have found, to our surprise, that the evidence is neither extensive nor particularly well known. The assumption of a bottom-to-top order of derivation seems to have arisen from some reasonable intuitions about human language. However, it is an assumption that has rarely been revisited over the years. Bottom-to-top is widely regarded as the only possible order of derivation, whereas it ought to be considered as just one among several possible ways to account for certain facts about language.

24

Probably the most influential intuition underlying this pervasive assumption is that lexical items are inserted into thematic positions. That is, a structure "begins" as a direct representation of thematic relations, which is subsequently transformed to satisfy other syntactic requirements. This intuition has its origins in the kernel sentences of *Syntactic Structures* (Chomsky, 1957), and it was codified in the Standard Theory as Deep Structure (later 'D-structure' or DS), a level of representation that "defines grammatical relations and functions in a straightforward manner" (Chomsky, 1970). It was the point in the derivation at which the argument selection requirements were satisfied (the 'Theta Criterion'). Empirically, DS was crucial to the distinction between raising and control. In raising constructions like (1), the subject bears a thematic relation to the embedded verb, but not the main verb. In control constructions like (2), the subject has a thematic relation to both verbs. It was claimed that raising constructions are formed by movement from a thematic position in the embedded clause to a non-thematic position in the main clause, whereas control constructions involve an empty category in the embedded clause that is not created by movement, but rather is created by *equi-NP deletion* or, in later parlance, the empty category PRO.

(1)   The man$_i$ seemed t$_i$ to enjoy his ice cream.

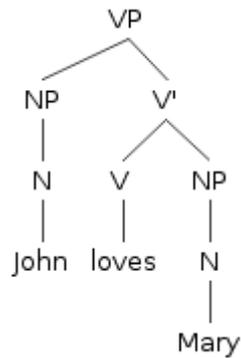(2)   The man$_i$ tried PRO$_i$ to enjoy his ice cream.

D-Structure was argued to be unnecessary as part of the Minimalist Program: the Theta Criterion was a requirement that might as well be satisfied at LF (Chomsky, 1993). Nevertheless, the new theory essentially reinstated it at the point of lexical insertion by requiring thematic roles to be assigned only at external Merge (Chomsky, 1995; but for a different view see Hornstein, 2000).

Applying the theta criterion at the point of lexical insertion is admittedly a fine way of capturing these data, and a bottom-to-top order of derivation follows from this. In derivations of this kind, satisfaction of the thematic requirements of a lexical item is guaranteed, and all further transformational operations serve to satisfy other syntactic requirements (encoded as "features"). However, this approach is by no means required. Stripping away theory-particular terminology, the aim is to account for the fact that arguments enter into multiple syntactic relations within a single structure. Usually only one of those relations (the one lowest in the structure) is thematic, except in the case of control constructions. These facts could just as well be captured by the opposite restriction: arguments must merge into a non-thematic position that satisfies syntactic requirements such as case, agreement or scope marking, and subsequent transformations could relate that argument to one or more lower positions in the structure where thematic and other requirements are satisfied. Arguably such operations are required under either framework: even in bottom-to-top theories, the original thematic relations must be re-established at LF for interpretation.
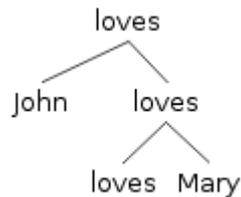
A second intuition underlying bottom-to-top derivation is that the endocentricity of syntactic structures is established by heads "projecting" phrases. The features of a head are said to be "passed up" to the maximal projection. Again, while such terminology is a convenient way of describing endocentricity, it is certainly not required. Bare Phrase Structure theory (cf. Chomsky, 1995) highlights the possibility that different levels of projection simply encode the relations between the head (a lexical item) and other parts of the structure. In the X'-style structure in (3), it might seem necessary to ensure that the features of the head 'travel' up the tree to give content to the otherwise empty categorial nodes. In the bare phrase structure in (4), the maximal projection of *loves* has no need to inherit features from the lexical item: it *is* that lexical item, represented in relation to other lexical items. In this approach, the labeling of levels of projection is merely a mechanism for identifying the dominant node, which determines how the phrase behaves as a whole. Endocentricity is a condition on the structure as a whole, and need not be explained by requiring the

projection of features from bottom to top. Mechanisms for regulating the match between the features of heads and maximal projections are familiar from constraint-based grammatical formalisms such as Head Driven Phrase Structure Grammar (HPSG), and can be straightforwardly implemented without the need for an ordered derivation.

(3)   X' Structure

```
              VP
            /    \
          NP      V'
          |      /  \
          N     V    NP
          |     |    |
        John  loves  N
                     |
                    Mary
```

(4)   Bare Phrase Structure

```
            loves
           /     \
        John     loves
                 /   \
             loves   Mary
```

We have argued that bottom-to-top derivation is not a logical necessity, despite the intuitive appeal of some long-standing assumptions. Still, since some version of this approach has been widely assumed for at least 40 years, we might expect that some other aspects of grammatical theory have been found to depend on it. In fact, our impression is that empirical arguments for this approach, to the extent that they are successful at all, are arguments for ordered derivations, but not arguments for a specific derivational order.

One aspect of many transformational grammatical theories that references the order of derivation is the syntactic cycle. In fact, Chomsky's (1966) formulation of cyclicity, which he refers to as a "general convention", is essentially a stipulation of bottom-to-top derivation. This restriction on the application of "singulary transformations" allowed the more powerful "generalized transformations" to be eliminated in favor of recursive base rules. The Strict Cycle Condition (SCC) (Chomsky, 1973) is slightly more sophisticated: it prohibits the operations in any particular syntactic domain from feeding operations in a subdomain. Some form of bottom-to-top derivational order follows from this requirement. A primary motivation for the SCC was to address Ross' (1967) observations about syntactic "islands", or barriers to long-distance movement. In a strictly bottom-to-top derivation, it is not possible to insert additional structure between a moved element and its original position after a licit movement is complete. That is, an element cannot escape being a barrier to a long-distance dependency simply by being added to the structure after the dependency is formed.

Freidin (1978) points out that though the empirical coverage of the SCC is substantial, it is subsumed by several other independently-motivated conditions on transformations. Freidin covers a number of different constructions; here we will mention just one example. The derivation in (5), resulting in the illicit sentence in (6), may be analyzed as a violation of the SCC.

26

(5)   a. [$_{CP}$ [$_{IP}$ John knows [$_{CP}$ [$_{IP}$ who$_i$ saw what$_j$]]]

     b. [$_{CP}$ [$_{IP}$ John knows [$_{CP}$ who$_i$ [$_{IP}$ t$_i$ saw what$_j$]]]

     c. [$_{CP}$ who$_i$ [$_{IP}$ John knows [$_{CP}$ t$_i$ [$_{IP}$ t$_i$ saw what$_j$]]]

     d. [$_{CP}$ who$_i$ [$_{IP}$ John knows [$_{CP}$ what$_j$ [$_{IP}$ t$_i$ saw t$_j$]]]

(6)   *Who does John know what saw?

In this derivation, the lower wh-phrase (*what*) moves after the higher wh-phrase (*what*) has already moved to the top of the tree. The movement in (5c) of a wh-phrase to the higher CP precedes the movement in (5d) of the other wh-phrase to the lower CP. The SCC rules out this derivation, since the lower CP is a cyclic subdomain of the higher CP. However, that step of the derivation also involves the deletion of a bound trace, which would presumably result in additional problems. Freidin invokes a "Trace Erasure Prohibition," which would rule out sentences like (6) regardless of the SCC. One could argue that Freidin's various conditions and principles are themselves stipulative, but his overall point is well taken, that cases that are excluded by the SCC are often already handled by independent constraints.

More generally, even if we allow that the examples in (5-6) motivate an ordered derivation in order to close off the theoretical loophole that might otherwise allow derivations like (5), such cases do not motivate a specific derivational order. Any restriction that forces derivations to proceed in a consistent order should do the trick. A derivation that proceeded from top-to-bottom would close the loophole in exactly the same way as does a bottom-to-top derivation. A strictly left-to-right derivation would equally block the derivation in (5), although it would conceivably make different predictions in situations where linear order and c-command relations diverge. Furthermore, examples like (5) are relevant to considerations of derivational order only due to a theory-internal assumption. If the locality constraint that rules out sentences like (6) is a constraint on movement operations (e.g. "wh-phrases mustn't move across other wh-phrases"), then there may be a need to block the derivation in (5). But if the locality constraint is a constraint on the representations created by movement (e.g., "a wh-phrase cannot form a dependency that crosses another wh-phrase"), then sentences like (6) will be ruled out irrespective of the order in which sentences are put together. This should come as no surprise: it is the reason why non-derivational grammatical theories are able to handle cases like (6).

Chomsky (1995) introduced the Extension Condition to capture the effects of cyclicity in a minimalist framework. According to the Extension Condition, all syntactic operations (i.e., Merge or Move) must target the root of the tree. Cyclicity follows immediately, as well as a prohibition on movement into complement positions. It has also been claimed that root Merge is "simpler" than other instances of Merge, and is thus preferable on Minimalist grounds. Along similar lines, Chomsky (2000, p. 136) argued that root-only merger is preferable as it has the effect that inserting new terminal nodes satisfies a 'least tampering' condition, because it does not destroy existing constituents.[8] Regardless of the status of these claims, however, Chomsky (1995) acknowledged that the Extension Condition could only apply before Spell-Out: covert movement operations seemed to be exempt. Adjuncts such as adverbials were also exempt. Subsequent formulations of cyclicity, such as Richards' (1999) "featural cyclicity," attempted to remedy the stipulative nature of the Extension Condition.

---

[8] Lasnik (2006) points out that this argument depends on the assumption of standard tree notation, and that it does not go through if other notational conventions are adopted.

However, they did not remedy the reliance of covert processes on non-root merger. If it is not the case that all syntactic operations can rely on root Merge alone, it seems arbitrary to require that pre-Spell-Out processes are restricted to it on the basis of its supposed simplicity. Moreover, notions of what counts as 'simpler' are subjective and perhaps not the most effective guide to figuring out the core properties of human language.

In the context of evidence for cyclic bottom-to-top derivations one often encounters discussion of successive cyclic movement. Successive cyclic movements are operations that displace a constituent across a long distance via a sequence of shorter steps. Successive cyclic movement was proposed as a way of reconciling evidence that on the one hand shows that phrases can be displaced across an arbitrarily long distance with evidence that on the other hand shows that many relatively local displacements are illicit. The proposal was that all displacement is subject to strict locality restrictions, and that apparent long-distance displacement is the result of a sequence of local operations (Chomsky, 1973). Some interesting cross-language phenomena have been offered as phonologically overt evidence of successive cyclic movement (e.g., Torrego, 1984; McCloskey, 1991). However, the properties of successive cyclic movement are compatible with a number of different derivational (or non-derivational) schemes. Successive cyclicity has attracted new interest in recent years in the context of minimalist notions of phase-based derivations (Chomsky, 2000; Boeckx, 2008). But again, we find little in these discussions that motivates a specific derivational order. Derivations that move elements from one phase to the next can be readily translated from a bottom-to-top order into other orders.

To summarize, we have argued that the evidence for strict bottom-to-top derivations is less than overwhelming. Some motivations for bottom-to-top derivations are the result of historical accidents, some reflect aesthetic choices about 'simplicity' that are likely to persuade only those who are already persuaded, and others may indeed motivate strictly ordering of syntactic operations, but do not entail that the specific order be bottom-to-top. Therefore we may pursue the idea that grammatical derivations reflect real-time processes with no concern that we are ignoring strong evidence for bottom-to-top derivations.

7.　Current status of arguments for (roughly) left-to-right syntactic derivations

Having spent much time discussing real-time structure building processes, and assessing the evidence for bottom-to-top syntactic derivations that look quite unlike those processes, we now turn our attention to the evidence for syntactic derivations that look somewhat more compatible with real-time mental computations.

Phillips (1996, 2003) argued that the constituency diagnostics that are the bread-and-butter of syntactic argumentation are better understood if we assume that sentence structures are assembled in a syntactic derivation that proceeds from left to right. This proposal was motivated entirely by the standard linguists' concern for distinguishing acceptable and unacceptable sentences. The possible psycholinguistic consequences were immediately apparent, however, and the proposal spawned an extensive experimental research program that has explored the nature of real time grammatical computation. We will not reiterate Phillips' arguments here in great detail, as the originals are readily available, but we will give an outline of the key proposal, in order to provide context for the responses that we discuss.

Phillips was concerned with a notorious problem in syntax. The syntactician's toolkit includes a variety of different diagnostics that can be used to probe the constituent structure of sentences. These diagnostics are based upon such processes as coordination, movement, ellipsis, substitution, scope, and binding, and they are a standard part of any introductory syntax course. In certain parade cases the diagnostics converge on the same result. For example, the diagnostics provide ample evidence for

28

the existence of a VP constituent: VPs can be coordinated, moved, elided, or substituted, and scope and binding relations show a clear subject-predicate asymmetry. However, a fact that is rarely discussed but well known to any practicing syntactician is that in many cases the diagnostics do not converge on the same results. Some diagnostics appear to apply to a broader range of constituents than others, some yield results that seem to be at odds with the results of others, and some even appear to contradict themselves by diagnosing conflicting constituent structures.

Coordination is often used as a diagnostic of constituenthood in introductory texts, but it is used much less often in syntactic research due to the fact that it is very liberal. In addition to stereotypical VPs (7a), coordination can target the two objects of double object constructions (7b), subject-verb combinations (7c), and sequences that span two clauses (7d).

(7)   a. Wallace [visited Wendolene] and [bought some wool].

       b. Wallace gave [Gromit a biscuit] and [Shawn some cheese] for breakfast.

       c. [Wallace designed] and [Gromit built] an enormous tin moon-rocket.

       d. Alice [knew Fred wanted to talk] and [hoped that he wanted to argue] with the president.

Other diagnostics are less liberal. For example, many strings that can be coordinated cannot be moved. The two objects of a double object construction can be coordinated, but they cannot be moved leftward or rightward in English (8). Similarly, ellipsis is more restrictive than coordination. Ellipsis can delete material from more than one clause (9a), but it cannot target a main clause plus a subpart of an embedded clause (9b), although the same sequence can be coordinated (5d).

(8)   a. *[Gromit a biscuit] Wallace gave ___ for breakfast.

       b. *Wallace gave ___ at breakfast time [his favorite pet beagle an enormous chewy biscuit.]

(9)   a. Alice [knew that Fred wanted to talk with the queen] and Ethel did ___ too.

       b. *Alice [knew that Fred wanted to talk with the queen] and Ethel did ___ with the president.

A pair of cases that have attracted a good deal of attention involve the relation between different structural diagnostics when they apply in the same sentence. The first case is a conflict noticed by Pesetsky (1995). The VP-fronting in (10), which places a verb and its two objects at the front of a clause and strands an adverbial PP, suggests a VP structure in which the predicate *give books to them* is a constituent and the adverbial PP attaches above this constituent. However, under the common assumption that binding is a diagnostic of c-command relations, then the successful licensing of the reciprocal *each other* suggests that the adverbial PP occupies a lower position where it is c-commanded by the pronoun *them*. These two structural diagnostics are thus in conflict with one another, and examples like (10) have come to be known as *Pesetsky's paradox*. A complementary case pointed out by Phillips (1996) involves VP-ellipsis, which targets very similar strings to VP-fronting, yet appears to avoid conflicts with c-command diagnostics involving binding and scope. For example, a pronoun inside an elided partial VP cannot bind a reciprocal inside a stranded adverbial, consistent with a structure in which the adverbial is structurally higher than the elided phrase (11). We will have more to say on this contrast below.

(10)   John wanted to give books to them$_i$, and [give books to them] he did ___ on each other$_i$'s birthdays. [VP-fronting]

(11)   *John wanted to give books to them$_i$ at Christmas, and Bill did ___ on each other$_i$'s birthdays. [VP-ellipsis]

The general problem of constituency conflicts has rarely been addressed directly. It has probably received the most attention in the Categorial Grammar (CG) literature, where the broad range of constituents identified by some diagnostics – coordination in particular -- has been offered as support for the CG mechanisms that allow words to be combined in different orders, leading to effects of flexible constituency (Steedman, 1985; Dowty, 1988; Pickering & Barry, 1993). Steedman (1997, 2000) offers a partial typology of structural diagnostics that accounts for why some tests yield different results than others. He argues that processes such as coordination, ellipsis, and unbounded dependencies are sensitive to the order of combinatorial operations in surface structure, and hence allow for effects of flexible constituency, whereas binding relations reflect prominence hierarchies in an independent semantic/argument structure representation that is insensitive to flexible combinatorial rules. Furthermore, Steedman's typology is not arbitrary, in the respect that processes that impact word order and surface form are associated with surface structure, whereas processes that impact interpretation are associated with the semantic structure representation. Although Phillips (2003) criticizes Steedman's approach on a number of grounds, arguing that it fails to predict which processes are (in)dependent of one another, the overall approach is highly appealing.

Pesetsky (1995) offers an account of constituency conflicts framed within a transformational grammar approach. Pesetsky proposes that different diagnostics yield contrasting results because sentences have more than one structural representation or derivation. He argues that every VP has a *cascade syntax* representation, which is largely right-branching, and a *layered syntax* representation, which is largely left-branching. Different structural diagnostics are assumed to be sensitive to one or the other of these representations, leading to an alternative typology of structural diagnostics. Setting aside specific empirical concerns about this approach (for discussion see Phillips, 2003), the main question that this approach raises involves the apparent arbitrariness of the typology.

Phillips (1996, 2003) offered a different approach to the typologies offered by Steedman and Pesetsky in an attempt to avoid the need to stipulate that individual structural diagnostics are sensitive to one type of representation and not another. He argued that the contrasting results of different structural diagnostics are natural consequences of the dynamics of left-to-right structure building. Starting from the assumption that humans have some kind of left-to-right structure building system for language, and assuming further that this system does not countenance vacuous projection of non-branching nodes,[9] some interesting predictions emerge about the evolution of constituent structure as a sentence is assembled. Whenever a word or phrase is merged with the current structure at any point other than the root node, some constituents cease to be constituents. For example, if the sentence *Harry met Sally* is assembled from left to right, then there is a stage in the derivation where *Harry met* is a constituent, consisting of the entire structure, but addition of the direct object *Sally*

---

[9] This assumption is not in conflict with experimental findings that suggest predictive creation of syntactic nodes before their phonological realization is encountered in the input. Predictive structure building in language comprehension demonstrates that structure-assembly is not fully synchronized with the words in the input. This is fully compatible with the assumption that all non-terminal nodes in syntax are branching nodes.

30

creates a new constituent *met Sally* while destroying the constituent *Harry met*. This means that the sequences of words that are constituents change over the course of a derivation. Phillips argues that the varying results of different constituency diagnostics are a direct consequence of this dynamic property of constituent structure, and he proposes that the results of individual diagnostics should be predictable, based on which stages of a left-to-right derivation they apply to. If a sequence of words is a constituent that does not survive to the end of the sentence, then it should only be visible to constituency diagnostics that can be applied before the constituent is destroyed.

Phillips presents a number of generalizations about structural diagnostics that he attributes to the dynamics of left-to-right derivations. For example, the liberal nature of coordination is argued to be a consequence of the fact that it places constituents adjacent to one another, and hence applies before constituents are destroyed. In contrast, ellipsis is a more restrictive diagnostic, because it involves structures in which an ellipsis site in a second clause (typically) is dependent on an antecedent in an earlier clause, and hence it can only 'see' the constituents that have survived to the end of the earlier clause. Meanwhile, structural diagnostics that test for c-command relations are predicted to not conflict with one another, since left-to-right expansion of a structure adds new c-command relations without destroying existing c-command relations.

Phillips also offers an account of why VP-fronting gives rise to Pesetsky's paradox, whereas VP-ellipsis does not. In cases of VP-fronting like (10) the movement operation applies before the stranded adverbial is added to the structure, and it is only the addition of the stranded adverbial (low inside the 'reconstructed' VP) that destroys the constituency of the fronted constituent. Therefore, tests of movement and binding appear to conflict with one another, because they apply at different stages in the assembly of the sentence. In cases of VP-ellipsis like (9), on the other hand, no corresponding conflict arises, because at the point when ellipsis must be licensed the right-branching VP-structure required to license the binding relation in (11) no longer has the sequence *give books to them* as a constituent. That sequence is certainly a constituent at some point during the construction of the first clause, but by the time the first clause is completed it is not a constituent, and hence it is 'invisible' for purposes of licensing ellipsis. As a result, the only structure that is compatible with the partial VP-ellipsis shown in (11) is one in which the stranded adverbial (and its counterpart in the first conjunct) attaches above the elided VP. This allows the relevant string to be a constituent at the point when ellipsis is licensed, but it is incompatible with the left-to-right binding relation.

The contrast between VP-fronting and VP-ellipsis was perhaps the most notable empirical result of Phillips' study, and it the only one to have attracted counter-arguments. Importantly, one counter-argument that Phillips anticipated was the possibility that the contrast in (10-11) might simply indicate that movement and ellipsis are subject to different constraints. Phillips addressed this concern by showing that the same contrast observed between VP-ellipsis and VP-fronting is also found within a single construction – comparative ellipsis. Comparative ellipsis shows the same structural properties as VP-fronting when it shares the word order properties of VP-fronting, and it shows the same structural properties as VP-ellipsis when it shares the word order of VP-ellipsis. This is illustrated in (12) using the scope of durational adverbials. (12a) exhibits an ambiguity between a collective reading, which describes the total number of books read within a single one week period, and a distributive reading, which describes the number of books that were each individually read in less than a week. Phillips assumes that this scope ambiguity reflects a structural difference between high attachment of the adverbial (collective) and low attachment below the direct object (distributive). Both of these attachment sites are compatible with the word order in (12a), in which the adverbial follows the ellipsis site and hence should not interfere with the licensing of ellipsis. On the other hand, only the high attachment

needed for the collective reading is compatible with (12b), in which an adverbial is in both the antecedent and comparative clauses, and hence interferes with the constituency required to license ellipsis. This behavior of comparative ellipsis makes it unlikely that the contrast in (10-11) merely reflects quirks of ellipsis and movement. Unfortunately, this argument from comparative ellipsis has been overlooked in all of the criticisms of the left-to-right account of the contrast in (10-11).

(12)　a. John read as many books as Bill did ___ in a week. [collective & distributive readings ok]

　　　b. John read as many books in a week as Bill did ___ in a month. [collective reading only]

The proposal that the results of different structural diagnostics reflect the dynamics of left-to-right structure building has much to recommend it. It addresses a fundamental problem in syntactic theory, it promises to bring together a diverse set of empirical facts under a general explanation, and it has interesting consequences for the mentalistic interpretation of grammatical theories. But it is important to ask whether the proposal has stood the test of time. The answer to this question appears to be a resounding "yes and no".

As a general account of why structural diagnostics yield different results the account has fared well, in the respect that no new accounts have emerged that offer a similarly broad account of the problem. A few studies have offered accounts of specific phenomena in terms of the more mainstream bottom-to-top derivations, and we will have more to say about those in a moment, but all of these have had a restricted empirical focus. There has been an upsurge of interest in coordination and right-node raising (e.g., Wilder, 1999; Sabbagh, 2007; Grosz, 2009; Larson, 2012); there is much new research on ellipsis (e.g., Merchant, 2001, 2004; Elbourne, 2008; Johnson, 2009); there is a rich set of findings on the interaction of scope and binding with movement processes (e.g., Fox, 1999; Boeckx, 2001); and so forth. But to our knowledge there have been no general attempts to explain why structural diagnostics yield diverse results.

However, one could argue that the general accounts of structural diagnostics (by Pesetsky, Phillips, Steedman, and others) have fared poorly in the respect that the issue that they highlight has not been widely acknowledged as a challenge for syntactic theory. We are unaware of efforts to tackle the broader question that was offered as evidence for left-to-right derivations. Additionally, the account of structural diagnostics in terms of left-to-right derivations has made little progress in terms of new empirical coverage. In particular, little work has tested whether the proposal succeeds or fails when extended to languages with different structural properties. This probably reflects the fact that there is little broad interest in the typology of structural diagnostics.

The proposal that structure building proceeds in a left-to-right fashion has also fared well in the respect that many new proposals have emerged that offer accounts of various linguistic phenomena in terms of left-to-right derivations. These include a number of proposals framed within a similar transformational phrase-structure grammar approach to Phillips, covering such phenomena as locality of movement (Richards, 1999; Chesi, 2007; Zwart, 2009), subject-auxiliary inversion (Bruening, 2008), reconstruction asymmetries (Guilliot, 2006), prosodic phrasing (Guimaraes, 1999; Wagner, 2005; Shiobara, 2008), expletive-associate relations (Richards, 1999), syntactic amalgams (Guimaraes, 2004), sprouting (Chung, Ladusaw, & McCloskey, 2011), right node raising (Park, 2005) and argument cluster coordination (Choi & Yoon, 2006), German 'SLF' coordination (Fortmann, 2005), and German adjective inflection (Schlenker, 1999). In addition, a number of interesting recent proposals have emerged in other grammatical frameworks in which left-to-right derivations play

32

a crucial role (e.g., Kempson, Gabbay, & Meyer-Viol, 2001; Cann, Kempson, & Marten, 2005; Shan & Barker, 2006; O'Grady, 2005).

On the other hand, proposals about left-to-right derivations have been relatively unsuccessful in the respect that there has been little genuine debate about the appropriateness of specific derivational orders, or indeed about the need for any derivations at all in syntax. Most views on these issues are currently relatively entrenched among syntacticians, and on all sides of the issue one encounters little interest in critically evaluating assumptions about how structures are put together. As a result the debate, such as it is, shows few signs of progressing towards a resolution.

We are aware of just a couple of proposals that aim to capture the generalizations that Phillips attributes to left-to-right derivations in terms of more standard bottom-to-top derivations. These have mostly focused on the account of the interaction between VP-fronting or VP-ellipsis with scope and binding (10-11). We will discuss these in a little more detail.

8. Scrambling in English (Baltin, 2003, 2006)

Baltin offers separate accounts of why partial VP-ellipsis bleeds binding possibilities (Baltin, 2003) and why partial VP-fronting does not (Baltin, 2006), within an approach that assumes bottom-to-top derivations and extensive scrambling of VP-internal phrases in English. Baltin argues that partial VP-ellipsis occurs when modifier phrases undergo string-vacuous rightward scrambling to a position above the VP, allowing the entire 'remnant' VP constituent to undergo deletion as a constituent. Given the additional assumption that binding constraints apply after scrambling, Baltin is able to capture the observation that material inside the elided VP cannot bind or take scope over elements inside the stranded modifier phrases (Baltin, 2003). This leaves Baltin with the challenge of explaining why fronting of partial VPs has a different effect than partial VP-ellipsis, such that elements inside the fronted VP are able to bind or take scope over stranded modifier phrases, as in (10). Previous analyses have taken it for granted that the binding relations must reflect a configuration in which the fronted VP is replaced ('reconstructed') in its underlying position, leading to Pesetsky's paradox. In contrast, Baltin assumes that the binding/scope relations reflect the surface position of the fronted VP. Of course, if the VP is fronted as a constituent, then the sub-constituents of the VP should fail to c-command elements outside of the VP. Instead, Baltin assumes that the fronting process that appears to move a VP constituent is in fact a series of movement operations that independently front the verb and scramble the other sub-constituents of the fronted VP (Baltin, 2006). The fact that these operations yield a word order that looks exactly like a fronted VP constituent is entirely coincidental under this account. By assuming that each of the sub-constituents of the VP is fronted independently, Baltin is able to explain how they are able to c-command modifier phrases that are stranded by the fronting process, thereby capturing the lack of interference between movement and binding. Thus, whereas Phillips argues that the contrast between partial VP-fronting and partial VP-ellipsis in (10-11) reflects the linear order properties of the two processes, Baltin proposes that they differ because neither is a process that really targets a partial VP constituent. Ellipsis applies to a larger full-VP constituent, which only appears to be smaller because of string-vacuous scrambling. Meanwhile, fronting applies to smaller units separately, only appearing to target a single larger constituent because the different fronting operations conspire to recreate exactly the same word order that existed prior to movement. Everything is not the way it seems.

Baltin's account of the interaction between VP-ellipsis and scope/binding (Baltin, 2003) has much in common with Phillips' account, in the respect that both accounts claim that ellipsis can strand VP-modifiers because the stranded modifiers occupy positions higher than the elided constituent. In both cases this captures the observation that a stranded anaphor cannot have an antecedent inside the elided material. The main

difference between the two accounts is that Baltin assumes that the stranded modifiers move from an initial low position to a higher position, which is actually to the left of the original VP constituent, whereas Phillips proposes that the stranded modifiers are initially merged in a high position. Baltin offers no reasons why his account of the English facts should be preferred. He does contend that his account better captures a parallelism between binding possibilities in English VP-ellipsis and Dutch scrambling, but since the Dutch facts simply show that binding possibilities respect surface word order, we do not find this consideration to be decisive. However, Baltin's article does present one simple generalization that does not follow from Phillips' account. Baltin points out that although binding relations are not possible between the elided and stranded VP material, ellipsis does not interfere with binding relations between multiple stranded modifiers (13).

(13) Tom played his guitar for the boys at their graduation, and Bob did so for the girls$_i$ on each other$_i$'s birthdays.

(13) illustrates the relatively mundane fact that the NP in the first stranded modifier *the girls* can bind a reciprocal in the second stranded modifier, just as it would if no ellipsis had occurred. This does not follow from Phillips' account, since he assumes that stranded modifiers are right-adjoined to VP, with the consequence that modifiers on the right should be structurally higher than those on the left. This (embarrassing) discrepancy calls for a modification of Phillips' account of the relative positioning of multiple stranded modifiers.

Baltin's account of the lack of interference between VP-fronting and scope/binding relations (Baltin, 2006) contrasts more strongly with Phillips' account of Pesetsky's paradox. Phillips adopts the standard assumption that VP-fronting involves fronting of a constituent, and that scope/binding relations are based on the non-fronted ('underlying') position of the VP. He attributes the apparent conflict between VP-fronting possibilities and scope/binding facts to the dynamics of left-to-right structure building. In contrast, Baltin assumes that Pesetsky's paradox arises in situations where a number of different movement operations proceed independently. First, the adverbials that appear to be stranded in sentences like (10) undergo scrambling out of the VP to a position where they ultimately appear to have not moved at all. Next, the remaining VP material is fronted as a constituent to a position below the subject NP (where none of that material ever surfaces). Following this, the arguments and modifiers that remain inside the VP are moved individually to positions above the subject NP, from where they c-command the rest of the clause to their right. Finally, the VP, which by this point contains only the verb, is moved to a Topic position at the front of the clause. This final step restores the linear order of the original VP constituent, although in Baltin's analysis it is not a constituent that has been fronted.[10]

Opinions may differ on how compelling one finds this way of capturing Pesetsky's paradox using standard bottom-to-top derivations. Importantly, though, Baltin offers empirical arguments that he claims favor his account over others. However, we think that these arguments do not ultimately favor his account.

---

[10] Baltin appears to assume that there are other circumstances where superficially similar word orders are derived by genuine fronting of a VP constituent, leading to cases where material inside the fronted VP cannot c-command stranded material to its right. Hence, some sentences may involve simpler derivations, but crucially these simpler derivations are not intended to account for cases of Pesetsky's paradox like (10).

34

First, Baltin offers a couple of arguments for analyzing the position of 'stranded' adverbials as being higher than the underlying position of the VP, contrasting with Pesetsky's and Phillips' assumption that the binding possibilities in (10) must reflect a right branching structure in which the stranded PP attaches low within the VP, such that it is c-commanded by the direct object NP. Baltin points to restrictions on British English *do* anaphora, a phenomenon that has much in common with the *do so* VP-anaphora found across many varieties of English. British English *do* anaphora, just like its *do so* counterpart, can strand true non-argument adverbials (14a), but cannot strand argument PPs (14b).

(14)   a. Although he wouldn't visit Sally$_i$ on her birthday, he would do ___ on her$_i$ anniversary.

       b. *Although I won't put the book on the table, I will do ___ on the mantelpiece.

Based on this and other properties of the construction, Baltin argues that the adverbials must be attached in a position higher than the rest of the VP. This is a reasonable analysis of the position of adverbials in VP-ellipsis, and in fact it is the same as the analysis of partial VP-ellipsis that Phillips (1996, 2003) adopts. However, since the challenge is to explain why the scope/binding properties of stranded adverbials differ between VP-fronting and VP-ellipsis, Baltin's argument from *do* anaphora, a form of VP-ellipsis, cannot be assumed to generalize to VP-fronting. In Phillips' account, it is the position of the stranded adverbial that differs between VP-fronting and VP-ellipsis, and in Baltin's account it is the position from which object NPs bind into the stranded adverbials that differs. The argument from *do* anaphora cannot decide between these two alternatives.

Baltin offers a further argument for the high position of adverbials stranded by VP-fronting, and this argument more clearly challenges the left-to-right account of Pesetsky's paradox. Baltin observes that examples like (15a), in which a quantificational direct object inside a fronted VP appears to license bound variable pronouns inside the stranded adverbial phrase, are exactly the kinds of phenomena that motivated Pesetsky's paradox. He points out that the variant in (15b), where the adverbial contains additional VP-ellipsis, creates a possible instance of antecedent contained deletion (ACD) if the adverbial is attached low inside the VP, as is assumed in Phillips' account. Phillips argued that low attachment of an adverbial inside a VP destroys the constituency of the verb and the direct object, making them invisible to subsequent syntactic processes. By this logic, the VP-ellipsis inside the adverbial in (15b) should not be able to take the VP *visit every prisoner* as its antecedent. Baltin argues that if the ellipsis is, in fact, possible, and if it does not interfere with establishing a bound variable interpretation for the pronoun, then this implies that the binding relations must reflect a higher position for the adverbial, and also a correspondingly higher position for the antecedent of the bound variable pronoun.

(15)   a. Visit every prisoner$_i$ though I may ___ after his$_i$ lawyer visits him$_i$, it won't matter.

       b. Visit every prisoner$_i$ though I may ___ after his$_i$ lawyer does ___, it won't matter.

Cases like (15b) certainly aim at the heart of what Phillips' account seeks to capture, by challenging the generalizations about the interactions between fronting, ellipsis, and scope/binding. However, we suspect that more work is needed to clarify the empirical facts in this domain before firm conclusions can be drawn. The status of (15b) and sentences like it strike us as unclear, and in related constructions the facts do not clearly favor Baltin's conclusion. The examples in (16) contain at most one bound variable pronoun, and offer the potential for a scope ambiguity between a collective reading, where the subject finished reading the set of books before Sally

did, and a distributive reading, where the subject read each individual book before Sally read that book. To the extent that (16b) is possible at all, the collective reading seems to be more readily available than the distributive reading, in contrast with (16a), where the singular pronoun forces the distributive reading. This suggests that VP-ellipsis may indeed change the structural position of the adverbial phrase, contrary to Baltin's argument. However, we should emphasize that Baltin's observation highlights an interesting domain where more work is needed.

(16)   a. Read every book$_i$ though he may before Sally reads it$_i$, she'll still get the better grades.

b. Read every book though he may before Sally does ___, she'll still get the better grades.

Baltin also argues that his account is supported by striking correlations between constraints on pseudogapping and constraints on Pesetsky paradox effects. According to his account, both pseudogapping and the partial VP-fronting that create Pesetsky paradox effects reflect scrambling operations that remove arguments from a VP. Thus, any non-trivial correlations between these two processes would provide striking support for his parallel analysis. Baltin's arguments are based on examples of pseudogapping and VP-fronting that involve complex VPs with more than one verb. (17) shows that complex predicates involving an infinitival under a subject control predicate allow both processes.

(17)   a. Although I didn't try to persuade Sally, I did ___ Martha.

b. Try to visit every prisoner$_i$ though I may ___ after his$_i$ lawyer does ___, I'm not sure they'll be successful.

In examples with more complex predicates pseudogapping is clearly degraded. (18a) contains a main clause object control verb, such that an overt NP appears between the two verbs, and the pseudogapped clause does not allow a construal in which the gap corresponds to *persuade Sally to visit*. Baltin argues that although the VP-fronting in (18b) is possible, it does not allow the indicated bound variable relation from the fronted direct object to the stranded adverbial, i.e., Pesetsky's paradox effects do not obtain. Similarly, Baltin points out that complex pseudogapping is impossible with predicates that select a non-finite complement headed by *from* (19a), and claims that when the same predicates undergo VP-fronting they also fail to show Pesetsky paradox effects (19b). Baltin argues that both of these restrictions are due to constraints on the scrambling processes that are part of his analysis of both constructions. Moreover, Baltin argues that the examples in (18c, 19c) that use coreferential pronouns instead of bound variable pronouns are acceptable, since coreference does not require c-command. He claims that these contrasts are mysterious under either Pesetsky's or Phillips' accounts.

(18)   a. *Although I didn't persuade Sally to visit Martha, I did ___ Susan.

b. *Persuade Sally to visit every student$_i$ though I may ___ on his$_i$ graduation day, it won't matter.

c. Persuade Sally to visit Tom$_i$ though I may ___ on his$_i$ graduation day, it won't matter.

(19)   a. *Although he didn't refrain from visiting Martha, he did ___ Susan.

b.*Refrain from visiting every student$_i$ though she may ___ on his$_i$ graduation day, it won't matter.

36

c. Refrain from visiting Tom_i though she may ___ on his_i graduation day, it won't matter.

These correlations are impressive if they are accurate, but we suspect that the restrictions are broader than Baltin suggests, and do not reflect constraints on scrambling. Baltin is probably correct that pseudogapping cannot target sequences that include an overt subject NP, as shown by the more minimal contrast in (20a-b). However, the degraded status of (20b) is probably not specifically due to the subject NP that intervenes between the main verb and the infinitival, since examples with intervening adverbials are similarly degraded (20c).

(20) a. Although I didn't want to visit Martha, I did ___ Susan.

        b. *Although I didn't want Sally to visit Martha, I did ___ Susan.

        c. *Although I didn't try as hard as I could to visit Martha, I did ___ Susan.

The contrasts in binding possibilities between (17b) on the one hand and (18b) and (19b) on the other hand are, frankly, too subtle for us to reliably assess. However, we suspect that the binding and coreference facts that Baltin points to are reflections of a more general constraint that affects VP-fronting and VP-ellipsis alike, and that was already discussed in Phillips (1998), an earlier version of Phillips (2003).[11] In any sequence consisting of *Aux [VP null] adverbial* the adverbial must be interpreted as a clausemate of the auxiliary. This can be seen clearly in the examples in (21), where the locative PP *in the basement* must be understood as modifying the matrix verb *resolve* rather than the embedded verb *fix*. The verb *resolve* is preferable to other control verbs like *want* and *try* used in some of the examples above, since the time and location of *want* and *try* and their complements coincide, a problem that does not arise with *resolve* and its complement.[12]

(21) a. Wallace resolved to fix the motorcycle in the garage, and Gromit did in the basement.

        b. Wallace needed to resolve to fix the motorcycle, and resolve to fix the motorcycle he did in the basement.

Phillips (1998) characterizes the clausemate constraint on modifier construal illustrated in (21) as an independent constraint that is unrelated to the dynamics of left-to-right structure building. Here we offer no deeper explanation of the constraint, although we note that it is not an unnatural constraint. We suggest that it is this constraint that is responsible for the contrasts that Baltin presents as a challenge for the left-to-right account.

Baltin's attempt to reconcile traditional bottom-to-top derivations with constituency conflicts faces a number of further challenges. First, Baltin's account of the contrast between VP-ellipsis and VP-fronting is based on construction-specific differences between the two phenomena, an approach that is challenged by the similar contrasts found internal to comparative ellipsis. Second, if Baltin is correct that VP-fronting is

---

[11] The 1998 manuscript includes some sections that were removed in the published version to satisfy length constraints.

[12] Landau (2007) shows that the clausemate constraint applies even with predicates like *want* and *try*, by using frequency adverbs, which disambiguate the main clause and embedded clause interpretations, as in (i).
(i) John intended to try to meet Mary often, but try to meet Mary he did only rarely. [rarely = try]

not a unitary operation, and that fronted arguments of the verb bind stranded modifiers from their fronted position, then we should expect that the same phrases should be able to bind subject NPs, a clearly inaccurate prediction (22). Third, Baltin's account of binding in VP-fronting is at odds with the well-supported independent generalization that fronted predicates obligatorily reconstruct for purposes of binding (Huang, 1993; Heycock, 1995; Takano, 1995). Fourth, it is puzzling under Baltin's account (as he acknowledges) that pseudogapping should allow argument stranding whereas VP-fronting does not (23). More generally, as argued by Landau (2007) partial VP-fronting is attested in many languages that lack productive scrambling rules (e.g., Italian, Spanish, Brazilian Portuguese, Hebrew), and partial VP-fronting in these languages allows stranding of phrases that cannot otherwise be removed from the VP. Finally, it is hard to overlook the fact that fronted VPs really do look like fronted constituents: they coordinate, they tolerate no interruptions, etc. This suggests that they really are constituents.

(22) a. *Visit every prisoner$_i$ though his$_i$ lawyer may ___ before the hearing, it won't matter.

     b. *Praise the girls$_i$ though each other$_i$'s mother did ___, they still felt sad.

(23) a. Bill visited Sally, and Bob ___ Martha.

     b.*Visit though he may ___ Sally, she won't be satisfied.

9. Remnant movement (Lechner 2003)

Lechner (2003) offers a bottom-to-top derivational account of the contrasting properties of VP-fronting and VP-ellipsis that, like Baltin's analysis, assumes that partial VP-fronting in English relies on counterparts of the scrambling and remnant VP-movement operations found in German. Lechner proposes that there are construction-specific differences in the way that modifiers are combined with the predicate in VP-fronting and VP-ellipsis constructions. In VP-fronting a modifier can be initially attached low in the VP, then later be removed from the VP by scrambling, clearing the way for partial VP-fronting, and can finally be reconstructed to its original site for purposes of evaluating binding relations. In VP-ellipsis, on the other hand, Lechner claims that the strategy of scrambling and then reconstructing the stranded modifier is not available, due to economy conditions that force the stranded modifier to directly merge in a position higher than the rest of the predicate. This high merger site accounts for why it is not possible for elided VP material to bind into the stranded modifier.

Since Lechner takes the position that there are construction-specific differences between partial VP-fronting and partial VP-ellipsis it is not clear how his account could capture the contrast within comparative ellipsis described above. A more detailed critique of Lechner's approach can be found in Landau (2007). More relevant to our current concerns, however, is to assess the empirical challenges that Lechner presents for the left-to-right structure building account of VP-fronting and VP-ellipsis.

Lechner points out that although partial VP-ellipsis blocks binding from an elided object into a stranded adverbial (9), no such restrictions are observed in the corresponding phrases in the first conjunct of a VP-ellipsis construction (24). This is unexpected under the assumption that VP-ellipsis must be licensed by a constituent in the antecedent clause. Similar examples were already discussed in Phillips (1996), where it was acknowledged that the judgments do not directly follow from the left-to-right structure building account.

(24) John gave the books to them$_i$ on each other$_i$'s birthdays, and Mary did ___ on their first day of school.

Lechner also points to scope freezing effects in partial VP-fronting that do not follow from the left-to-right account. In (25) the stranded PP obligatorily takes wide scope relative to the fronted predicate, although Phillips' account predicts that it should be possible to merge it in a low position inside the predicate, once the predicate is copied to its thematic position.

(25)   a. David planned to deliver every book to one of the students,$\forall > \exists, \exists > \forall$

      b. …[$_{VP}$ and deliver every book] he did ___ to one of the students.*$\forall > \exists, \exists > \forall$

More can be said about the details of these observations (see Phillips, 2002; Lechner, 2002 for discussion), but we suspect that the ultimate source of these effects lies in incremental interpretation. Once the binding relation in the first clause of (24) is established, it cannot be retracted based on the constituency requirements of the VP-ellipsis construction. The fronted predicate in (25) is interpreted in its fronted position, such that subsequent phrases must take scope outside this interpretive unit. These suggestions are certainly speculative, and they beg a number of questions, but we suspect that it is considerations such as these, rather than raw constituency constraints, that are ultimately responsible for the unexpected interpretations.

Lechner's third criticism of the left-to-right approach parallels one of Baltin's concerns. He points to examples of antecedent contained deletion (ACD) in a phrase stranded by VP-fronting (26), and argues that these should be impossible under the left-to-right structure building analysis. The problem is that if the stranded phrase merges low inside the fronted predicate, then the antecedent for the ellipsis site contains the ellipsis site, leading to an infinite regress. It is important that these cases involve stranded (but optional) argument phrases, as these are not assumed to have the option of simply merging high in the VP to avoid the constituency problem.

(26)   Mary asked him to deliver a book to some of the boys, and [$_{VP}$ deliver a book] he did ___ to every boy Mary wanted him to ___ .

The second ellipsis site in (26) is interpreted roughly as *deliver a book to x*, where the variable is bound by the quantificational NP *every boy*. Lechner is correct that the availability of this antecedent for ellipsis does not follow directly from the left-to-right structure building account. In fact, the puzzle that such cases present can be seen in simpler examples that lack VP-fronting, such as (27).

(27)   Harry delivered a book to every boy Mary wanted him to ___ .

We should emphasize that the puzzle that such examples present appears to be specific to ACD. In other instances of VP-ellipsis, judgments are as expected. The examples in (28) show VP-ellipsis inside a relative clause, as in ACD, but the head of the relative clause is non-quantificational, and the ellipsis site is marked by *do so*. The ellipsis is fine in (28a), where the ellipsis antecedent *read a book* is a constituent that does not include the ellipsis site. The ellipsis is impossible in (28b), where the intended ellipsis antecedent contains the ellipsis site. Since regular VP-ellipsis behaves as expected, and ACD in examples like (26-27) creates exactly the same constituency puzzle for the left-to-right approach that it creates for traditional approaches, we assume that cases like (26-27) should be handled in the left-to-right approach using similar mechanisms to those used to handle ACD in other approaches (e.g., quantifier raising).

(28)   a. John read a book at a table where Mary had done so ___ half an hour earlier.

      b. *John put a book on a table where Mary had done so ___ half an hour earlier.

Lechner's final concern involves the clausemate constraint on the construal of phrases stranded by VP-ellipsis and VP-fronting. He notes that the constraint amounts to a restatement of the facts, a criticism that is largely accurate. There is nothing inherently wrong with the fact that the generalization is not a direct consequence of left-to-right structure building, since we should not expect derivational order to be a cure-all that explains all generalizations. But it is, of course, interesting if the generalization turns out to follow from independent principles. Lechner offers one such proposal, which essentially treats the clausemate constraint as a consequence of locality restrictions on extraposition processes. Landau (2007) provides empirical arguments against Lechner's interpretation of the restriction, and proposes an alternative account of his own. In our impression, the jury remains out on the appropriate analysis of the clausemate constraint.

10. Late merger of adjuncts (Landau 2007)

Landau (2007) offers an interesting discussion of constraints on partial VP-fronting, with a focus on the requirement that fronted VPs be at least partially complete VP constituents (Phillips, 2003). Landau discusses data from English and from Hebrew, a language that appears to be particularly well suited to examining partial VP-fronting. His account of Pesetsky's paradox is framed in terms of bottom-to-top derivations, but his account has much in common with the left-to-right structure building account. Landau resolves the paradox by assuming that VP-fronting occurs prior to merger of stranded adjunct phrases into the null VP-copy. Thus the fronted VP is a constituent at the point where movement occurs, although it is no longer a constituent in its underlying position once the stranded adjunct is inserted into the structure. This order-based solution is very similar to Phillips' account, although the two proposals differ in a number of details.

Landau criticizes the left-to-right structure building account on the grounds that it does not offer a compelling account of the requirement that fronted VPs be potentially complete VPs. Phillips (2003) did offer a half-hearted explanation, and Landau does a good job of demonstrating the limitations of that account. Landau offers his own account of the restriction, based on a requirement that verbs' thematic requirements be satisfied locally. If Landau is correct, then his account of the constraint could straightforwardly be implemented in a wide variety of grammar formalisms, independent of derivational order, and so it could be readily incorporated into the left-to-right account.

Landau offers argues that his approach is supported by an interesting additional generalization, based on an extension of well-known asymmetries in reconstruction effects pointed out by Freidin (1986) and Lebeaux (1988). Freidin and Lebeaux pointed out that clausal arguments of fronted NPs show reconstruction effects for purposes of Condition C (29a, 30a), but relative clauses do not. If the entire fronted NP in (29b, 30b) is interpreted in direct object position, then the indicated coreference relation should be ruled out by Condition C, yet coreference is possible.

(29)   a. ?*Which argument that John$_i$ is a genius did he$_i$ believe?

   b. Which argument that John$_i$ made did he$_i$ believe?(Lebeaux, 1988)

(30)   a. ?*The overwhelming evidence that Henry$_i$ was a spy, he$_i$ refused to accept.

   b. The overwhelming evidence that Henry$_i$ had amassed, he$_i$ refused to present.

A popular approach to this paradigm is to assume that relative clauses and other adjuncts are able to escape the effects of Condition C because they are merged with the wh-phrase in its fronted position, and hence are never c-commanded by the coreferential pronoun.[13] Landau offers an interesting twist on this generalization, arguing that there is a correlation between the phrases that can escape Condition C effects when they are fronted and the phrases that can be stranded when a predicate is fronted. This correlation is expected under Landau's approach, since both effects are attributed to late merger of the non-obligatory phrases. Landau extends his account to cover an additional contrast between different types of modifiers inside fronted predicates. Although a relative clause inside a fronted nominal is able to escape Condition C effects (31a), the Condition C effect reemerges when the same nominal is embedded inside a fronted predicate (31b; Takano, 1995). In contrast, Landau observes that Condition C can be avoided when the critical noun is part of a VP-modifier in a fronted predicate (31c). Landau argues that this can be explained by his account, if late adjunction of adjuncts is further constrained such that it cannot apply *inside* predicates. The modifier containing the noun *Mary* is inside the fronted predicate in (31b) but is presumably adjoined to the top of the fronted predicate in (31c). Landau suggests that the constraint on late adjunction inside predicates could also explain the clausemate condition on the interpretation of stranded modifiers.

(31)   a. Food that Mary$_i$ cooks, she$_i$ knows I would never eat.

     b.*Eat food that Mary$_i$ cooks, she$_i$ knows I never would.

     c. Eat food at Mary$_i$'s party, she$_i$ knows I never would.

Interestingly, the contrasting possibilities for escaping Condition C effects may be equally compatible with a left-to-right structure building account, without a need to invoke late adjunction processes. Phillips (1998; section removed in 2003 version) offers an account of the Freidin-Lebeaux facts in (29-30), proposing that the contrast simply reflects constituent structure. If it is assumed that a relative clause is adjoined to NP, then there is a smaller constituent (e.g., *which argument*) that may undergo reconstruction, without reconstructing the critical name and inducing a Condition C violation. On the other hand, if clausal complements of nouns are sisters of N, then there is no constituent that includes the *wh*-phrase *which argument* and excludes the complement clause, and consequently the entire nominal must reconstruct, leading to a Condition C violation. This account straightforwardly extends to the contrast in (31ab). (31a) is structurally identical to (29b, 30b), and hence the relative clause may avoid reconstruction. But once the nominal is embedded inside a predicate (31b) there is no constituent that includes the verb and the direct object and excludes the relative clause, and thus the entire predicate must reconstruct.

---

[13] Several authors have questioned the accuracy of the Freidin-Lebeaux contrast (e.g., Bianchi, 1995; Kuno, 1997; Postal, 1997; Lasnik, 2003). These authors have pointed to examples like (i-iii), which suggest that extraction of arguments of NP may sometimes also bleed Condition C.

(i)Which piece of evidence that John$_i$ was guilty did he$_i$ successfully refute? (Lasnik, 2003)
(ii)The claim that the director$_i$ was corrupt, he$_i$ was unwilling to discuss. (Postal, 1997)
(iii)Whose claim that the senator$_i$ had violated the campaign finance regulations did he$_i$ dismiss as politically motivated? (Kuno, 1997)

The facts in (i-iii) are interesting, but we do not think that they undermine the interest of the Freidin-Lebeaux contrast. Adjuncts to NP appear to reliably escape Condition C effects, whereas there is some variability regarding the status of arguments of NP. Landau (2007) suggests that this may be because some finite complements that have been analyzed as nominal complements should, in fact, be treated as adjuncts (following Stowell, 1981; Safir, 1999).

The status of reconstruction for VP-modifiers is less clear. Landau argues based on examples like (31c) that such modifiers can freely escape Condition C effects. But others, going back as far as Hasegawa (1983; observation attributed to Joan Bresnan) have argued that fronting of predicates can never bleed Condition C. Phillips (1998) offers the examples in (32) of VP-modifiers, which appear to be subject to obligatory reconstruction.

(32) a. *[Playing cards until long after Lucy$_i$'s bedtime] though she$_i$ was, she$_i$ was not at all tired the next morning.

      b. *[Read a biography in Wallace$_i$'s living room] though he$_i$ would like to, there's no chance that he$_i$ actually will.

The modifiers that undergo obligatory reconstruction in (32) are modifiers that would normally allow stranding in cases of VP-fronting (33). Such cases cast doubt upon Landau's correlation between modifiers that escape Condition C and that can be stranded by predicate fronting. More generally, though, the contrast between (31c) and (32) suggests that more work is needed to understand under what circumstances VP-modifiers can escape Condition C effects.

(33) a. Play cards though Lucy did ___ until long after her bedtime, she was not at all tired the next morning.

      b. Read a biography though Wallace did ___ in his living room, he still couldn't relax.

Summarizing, a number of specific challenges have been raised for the account of constituency conflicts offered by Phillips (1996, 2003), framed in terms of more 'conservative' approaches to syntactic derivations, i.e., derivations that start at the end of a sentence and end at the beginning. Some of these challenges can be addressed easily, while others are less straightforward. In some cases, such as Landau's account of the potential complete VP constraint, interesting new generalizations could be straightforwardly grafted onto the left-to-right account. We are unaware of any alternative proposals that have the scope of the left-to-right account, or that offer such a compelling bottom-to-top analysis to motivate abandoning the left-to-right account.

11. Conclusion

Discussions of left-to-right order in syntactic derivations tend to focus on standard considerations of descriptive adequacy – whether left-to-right order provides broader coverage of (un)acceptable sentences. But the elephant in the room in such discussions is the possibility of understanding human grammatical computation in more realistic terms, as mental processes that operate in real time and that are open to scrutiny using a host of cognitive and neuroscientific tools. Here we have argued that current evidence favors the view that human grammatical abilities are best understood as a single structure building system that assembles syntactic structures in a (roughly) left-to-right order, and is a key component of systems for parsing and production. This position departs from widespread assumptions in linguistics in two ways. For at least 40 years it has been standard to assume that human grammatical competence is an implementation independent system, and that this system assembles sentences via bottom-to-top derivations that typically proceed in the opposite order than normal comprehension and production. We have argued that there is very little evidence for the implementation independence of human grammatical abilities, and that the motivations for bottom-to-top syntactic derivations are less than overwhelming. Meanwhile, the feasibility of the real-time approach to grammatical computation has been reinforced by linguistic and psycholinguistic evidence of recent years. This implies a research program for grammatical theory that goes far beyond the traditional concern with classifying acceptable and unacceptable sentences, and it suggests that

42

grammatical theories should be accountable to a much richer body of evidence, particularly evidence on the time course of grammatical processes. This is just what we should want.

References

Aoshima, S., Phillips, C., & Weinberg, A. (2004) Processing filler-gap dependencies in a head-final language. *Journal of Memory and Language, 51*, 23-54.

Aoshima, S., Yoshida, M., & Phillips. C. (2009) Incremental processing of coreference and binding in Japanese. *Syntax, 12*, 93-134.

Badecker, W. & Straub, K. (2002) The processing role of structural constraints on the interpretation of pronouns and anaphora. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 748-769.

Baltin, M. (2003) The interaction of ellipsis and binding: Implications for the sequencing of Principle A. *Natural Language and Linguistic Theory, 21*, 215-246.

Baltin, M. (2006) The non-unity of VP preposing. *Language, 82*, 734-766.

Bever, T. G. (1970) The cognitive basis for linguistic structures. In R. Hayes (ed.), *Cognition and Language Development*, pp. 279-362. New York: Wiley.

Bever, T. G., Lackner, J. R., & Kirk, R. (1969) The underlying structures of sentences are the primary units of immediate speech processing. *Perception and Psychophysics, 6*, 225-234.

Bianchi, V. (1995) *Consequences of antisymmetry for the syntax of headed relative clauses*. PhD dissertation, Scuola Normale Superiore, Pisa.

Bock, J. K. (1986) Syntactic persistence in language production. *Cognitive Psychology, 18*, 355-387.

Bock, J. K. (1989) Closed-class immanence in sentence production. *Cognition, 31*, 163-186.

Bock, J. K. & Loebell, H. (1990) Framing sentences. *Cognition, 35*, 1-39.

Boeckx, C. (2001) Scope reconstruction and A-movement. *Natural Language and Linguistic Theory, 19*, 503-548.

Boeckx, C. (2008) *Understanding minimalist syntax: Lessons from locality in long-distance dependencies*, Oxford, UK: Blackwell.

Bruening, B. (2008) Locative inversion and *do*-support. Ms., University of Delaware.

Brouwer, H., Fitz, H., & Hoeks, J. (2012) Getting real about semantic illusions: rethinking the functional role of the P600 in language comprehension. *Brain Research, 1446*, 127-143.

Cann, R., Kempson, R., & Marten, L. (2005) *The dynamics of language.* San Diego: Academic Press.

Carnie, A. (2006) *Syntax: A generative introduction, 2nd edition.* Malden, MA: Wiley-Blackwell.

Chesi, C. (2007) Five reasons for building phrase structures top-down from left to right. *Nanzan Linguistics: Special Issue 3*, 71-105.

Choi, Y. & Yoon, J. (2006) Argument cluster coordination and constituency test (non-)conflicts. Talk at *NELS 37*, University of Illinois, Urbana-Champaign.

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

Chomsky, N. (1965) *Aspects of the theory of syntax*, Cambridge, MA: MIT Press.

Chomsky, N. (1966) *Topics in the theory of generative grammar*, The Hague: Mouton.

Chomsky, N. (1970) Deep structure, surface structure, and semantic interpretation. In R. Jakobson and S. Kawamoto (eds.), *Studies in General and Oriental Linguistics*, Tokyo, Japan: TEC Company, Ltd., 52-91.

Chomsky, N. (1973) Conditions on transformations. In S.R. Anderson and P. Kiparski (eds.), *A Festschrift for Morris Halle*, New York, NY: Holt, Rinehart, and Winston, 232-86.

Chomsky, N. (1981) *Lectures on government and binding*. Dordrecht: Foris.

Chomsky, N. (1993) A minimalist program for linguistic theory. In K. Hale and S.J. Keyser (eds.), *The View from Building 20: Essays in Linguistics in Honor of Sylvain Bromberger*, Cambridge, MA: MIT Press, 1-52.

Chomsky, N. (1995) *The Minimalist Program*, Cambridge, MA: MIT Press.

Chomsky, N. (2000) Minimalist inquiries: The framework. In R. Martin, D. Michaels, and J. Uriagereka (eds.), *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*, Cambridge, MA: MIT Press, 89-155.

Christianson, K., Hollingworth, A., Halliwell, J., & Ferreira, F. (2001) Thematic roles assigned along the garden path linger. *Cognitive Psychology, 42*, 368-407.

Chung, S., Ladusaw, W., & McCloskey, J. (2011) Sluicing: between structure and inference. In R. Gutiérrez-Bravo, L. Mikkelsen, & E. Potsdam (eds.), *Representing language: Essays in honor of Judith Aissen*, pp. 31-50. Santa Cruz, CA: Linguistic Research Center.

Clifton, C. Jr., Frazier, L., & Deevy, P. (1999) Feature manipulation in sentence comprehension. *Rivista di Linguistica, 11*, 11-39.

Clifton, C., Jr., Kennison, S., & Albrecht, J. (1997) Reading the words *her*, *him*, and *his*: Implications for parsing principles based on frequency and on structure. *Journal of Memory and Language, 36*, 276-292.

Culicover, P. & Jackendoff, R. (2005) *Simpler syntax*. Oxford, UK: Oxford University Press.

deLong, K. A., Urbach, T. P., & Kutas, M. (2005) Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience, 8*, 1117-1121.

Dillon, B., Mishler, A., Sloggett, S., & Phillips, C. (in press) Contrasting interference profiles for agreement and anaphora. *Journal of Memory and Language.*

Dowty, D. (1988) Type raising, function composition, and non-constituent conjunction. In R. T. Oehrle, E. Bach, and D. Wheeler (eds.), *Categorial grammars and natural language structures*, pp. 153-197. Dordrecht: Reidel.

Drenhaus, H., Saddy, D., & Frisch, S. (2005) Processing negative polarity items: When negation comes through the back door. In S. Kepser & M. Reis (eds.), *Linguistic evidence – empirical, theoretical, and computational perspectives*, pp. 145-165. Berlin/New York: Mouton de Gruyter.

Elbourne, P. (2008) Ellipsis sites as definite descriptions. *Linguistic Inquiry, 39*, 191-220.

Federmeier, K. & Laszlo, S. (2009) Time for meaning: electrophysiology provides insights into the dynamics of representation and processing in semantic memory. In B. Ross (ed.), *The Psychology of Learning and Motivation, vol. 51*, pp. 1-44. Burlington, VT: Academic Press.

Ferreira, F. (2003) The misinterpretation of noncanonical sentences. *Cognitive Psychology, 47*, 164-203.

Ferreira, F. & Patson, N. D. (2007) The 'good enough' approach to language comprehension. *Language and Linguistics Compass, 1*, 71-83.

Fillenbaum, S. (1971) Psycholinguistics. *Annual Review of Psychology, 22*, 251-308.

Fodor, J. A., Bever, T. G., & Garrett, M. F. (1974). *The psychology of language*. New York: McGraw-Hill.

Fortmann, C. (2005) Die Lücken im Bild von der Subjectlücken-Koordination. *Linguistische Berichte, 204*, 441-476.

Fox, D. (1999) Reconstruction, binding theory, and the interpretation of chains. *Linguistic Inquiry, 30*, 157-196.

Fox Tree, J. E. & Meijer, P. J. A. (1999) Building syntactic structure in speaking. *Journal of Psycholinguistic Research, 28*, 71-90.

Frazier, L. (1987) Sentence processing: A tutorial review. In M. Coltheart (ed.), *The psychology of reading*, pp. 559-586. Hove, UK: Erlbaum.

Freidin, R. (1978) Cyclicity and the theory of grammar. *Linguistic Inquiry* 9, 519-549.

Freidin, R. (1986) Fundamental issues in the theory of binding. In B. Lust (ed.), *Studies in the acquisition of anaphora,* pp. 151-188. Dordrecht: Reidel.

Friederici, A. D., Pfeifer, E., & Hahne, A. (1993) Event-related brain potentials during natural speech processing: Effects of semantic, morphological and syntactic violations. *Cognitive Brain Research, 1*, 183-192.

Friederici, A. D. & Weissenborn, J. (2007) Mapping sentence form onto meaning: The syntax-semantic interface. *Brain Research, 1146*, 50-58.

Gouvea, A. C., Phillips, C., Kazanina, N., & Poeppel, D. (2010) The linguistic processes underlying the P600. *Language and Cognitive Processes, 25*, 149-188.

Grodner, D., Gibson, E., Argaman, V., & Babyonyshev, M. (2003) Against repair-based reanalysis in sentence comprehension. *Journal of Psycholinguistic Research, 32,* 141-166.

Grosz, P. (2009) Movement and agreement in right-node raising constructions. Talk presented at GLOW 32, Nantes, France.

Guilliot, N. (2006) A top-down analysis for reconstruction. *Lingua, 116*, 1888-1914.

Guimaraes, M. (1999) Deriving prosodic structure from dynamic top-down syntax. Ms., University of Maryland.

Guimaraes, M. (2004) *Derivation and representation of syntactic amalgams*. PhD dissertation, University of Maryland.

Hagoort, P. (2003) How the brain solves the binding problem for language: A neurocomputational model of syntactic processing. *NeuroImage, 20*, S18-S29.

Hagoort, P., Brown, C. M., & Groothusen, J. (1993) The Syntactic Positive Shift (SPS) as an ERP measure of syntactic processing. *Language and Cognitive Processes, 8*, 439-483.

Harris, C. L. & Bates, E. A. (2002) Clausal backgrounding and pronominal reference: A functionalist approach to c-command. *Language and Cognitive Processes, 17*, 237-269.

Hawkins, J. (1994) *A performance theory of order and constituency*. Cambridge, UK: Cambridge University Press.

Heycock, C. (1995) Asymmetries in reconstruction. *Linguistic Inquiry, 26*, 547-570.

Hofmeister, P. & Sag, I. A. (2010) Cognitive constraints on syntactic islands. *Language, 86*, 366-415.

Hornstein, N. (2000) *Move: A minimalist theory of construal*. Oxford: Blackwell.

44

Huang, C.-T. J. (1993) Reconstruction and the structure of VP: Some theoretical consequences. *Linguistic Inquiry, 24*, 103-138.

Inoue, A. & Fodor, J. D. (1995) Information-paced parsing of Japanese. In R. Mazuka & N. Nagai (eds.), *Japanese sentence processing*, pp. 9-63. Hillsdale, NJ: Erlbaum.

Jackendoff, R. (2002) *Foundations of language*. Oxford: Oxford University Press.

Johnson, K. (2009) Gapping is not (VP-)ellipsis. *Linguistic Inquiry, 40, 289-328.*

Kazanina, N., Lau, E. F., Yoshida, M., Lieberman, M., & Phillips, C. (2007) The effect of syntactic constraints on the processing of backwards anaphora. *Journal of Memory and Language, 56*, 384-409.

Kempson, R., Meyer-Viol, W., & Gabbay, D. (2001) *Dynamic syntax: The flow of language understanding*. Malden, MA: Blackwell.

Kennison, S. (2003) Comprehending the pronouns *her*, *him*, and *his*: Implications for theories of referential processing. *Journal of Memory and Language, 52*, 205-225.

Kluender, R. & Kutas, M. (1993) Subjacency as a processing phenomenon. *Language and Cognitive Processes, 8*, 573-633.

Kuno, S. (1997) Binding theory and the minimalist program. Ms. Harvard University.

Landau, I. (2007) Constraints on partial VP-fronting. *Syntax, 10*, 127-164.

Larson, B. (2012) A dilemma with accounts of right node raising. *Linguistic Inquiry, 43*, 143-150.

Lasnik, H. (2003) *Minimalist investigations in linguistic theory*. London: Routledge.

Lasnik, H. (2006) Conceptions of the cycle. In L. Cheng & N. Corver (eds.), *Wh-movement moving on*, pp. 197-216. Cambridge, MA: MIT Press.

Lau, E. F. (2009) *The predictive nature of language comprehension*. PhD dissertation, University of Maryland.

Lau, E. F., Phillips, C., & Poeppel, D. (2008) A cortical network for semantics: (de)constructing the N400. *Nature Reviews Neuroscience, 9*, 920-933.

Lau, E. F., Stroud, C., Plesch, S., & Phillips, C. (2006) The role of structural prediction in rapid syntactic analysis. *Brain and Language, 98*, 74-88.

Lebeaux, D. (1988) *Language acquisition and the form of the grammar*. PhD dissertation, University of Massachusetts.

Lechner, W. (2002) Some comments on Phillips (2002). Ms., Institute of General and Theoretical Linguistics, University of Tübingen.

Lechner, W. (2003) Phrase structure paradoxes, movement, and ellipsis. In K. Schwabe & S. Winkler (eds.), *Adding and omitting*, pp. 177-203. Amsterdam, John Benjamins.

Lee, M.-W. & Williams, J. (2006) The role of grammatical constraints in intra-sentential pronoun resolution. Manuscript, London Metropolitan University and University of Cambridge.

Legendre, G., Grimshaw, J., & Vikner, S. (2001) *Optimality-theoretic syntax*. Cambridge, MA: MIT Press.

Levelt, W. J. M. (1970) Hierarchical chunking in sentence processing. *Perception and Psychophysics, 8*, 99-103.

Levelt, W. J. M. (1974) *Formal grammars in linguistics and psycholinguistics* (3 volumes). The Hague: Mouton.

Lewis, S., Chow, W.-Y., & Phillips, C. (2012) Structural constraints on pronoun resolution: distinguishing early and late sensitivity to illicit antecedents. Talk at GLOW 35, Potsdam, Germany.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994) The lexical nature of syntactic ambiguity resolution. *Psychological Review, 101*, 676-703.

Marr, D. (1982) *Vision*. San Francisco: W. H. Freeman.

Mazuka, R. & Itoh, K. (1995) Can Japanese speakers be led down the garden path? In R. Mazuka & N. Nagai (eds.), *Japanese sentence processing*, pp. 295-329. Hillsdale, NJ: Lawrence Erlbaum Associates.

McCloskey, J. (1991) Resumptive pronouns, A' biding, and levels of representation in Irish. In R. Hendrick (ed.), *Syntax and Semantics 23: The syntax of the modern Celtic languages*, pp. 199-248. New York: Academic Press.

Merchant, J. (2001) *The syntax of silence: sluicing, islands, and the theory of ellipsis.* Oxford University Press.

Merchant, J. (2004) Fragments and ellipsis. *Linguistics and Philosophy, 27*, 661-738.

Momma, S., Slevc, R., & Phillips, C. (2013) Advance planning of verbs in head-final language production. Poster at the 26[th] annual CUNY Conference on Human Sentence Processing, Columbia, S. Carolina.

Montalbetti, M. (1984) *After binding*. PhD dissertation, Massachusetts Institute of Technology.

Neeleman, A., & van de Koot, H. (2010). Theoretical validity and psychological reality of the grammatical code. In H. De Mulder, M. Everaert, Ø. Nilsen, T. Lentz, & A. Zondervan (eds.), *Theoretical validity and psychological reality*, pp. 183-212. Amsterdam: John Benjamins.

Neville, H. J., Nicol, J., Barss, A., Forster, K. I., & Garrett, M. F. (1991) Syntactically based sentence processing classes: Evidence from event-related brain potentials. *Journal of Cognitive Neuroscience, 3*, 151-165.

Nicol, J. & Swinney, D. (1989) The role of structure in coreference assignment during sentence comprehension. *Journal of Psycholinguistic Research, 18*, 5-19.

O'Grady, W. (2005) *Syntactic carpentry: An emergentist approach to syntax*. Mahwah, NJ: Erlbaum.

Osterhout, L. & Holcomb, P. J. (1992) Event-related potentials elicited by syntactic anomaly. *Journal of Memory and Language, 31,* 785-806.

Park, M.-K. (2005) When things are cumulated or distributed across coordinate conjuncts. *Studies in Generative Grammar, 15*, 415-431.

Pearlmutter, N. J., Garnsey, S. M., & Bock, J. K. (1999) Agreement processes in sentence comprehension. *Journal of Memory and Language, 41*, 427-456.

Pesetsky, D. (1995) *Zero syntax*. Cambridge, MA: MIT Press.

Phillips, C. (1996) *Order and structure*. PhD dissertation, Massachusetts Institute of Technology.

Phillips, C. (2002) Constituency in deletion and movement. Ms., University of Maryland.

Phillips, C. (2003) Linear order and constituency. *Linguistic Inquiry, 34*, 37-90.

Phillips, C. (2004) Linguistics and linking problems. In M. Rice & S. Warren (eds.), *Developmental language disorders: From phenotypes to etiologies*, pp. 241-287. Mahwah, NJ: Erlbaum.

Phillips, C. (2006) The real-time status of island phenomena. *Language, 82*, 795-823.

Phillips, C. (2013a) Some arguments and non-arguments for reductionist accounts of syntactic phenomena. *Language and Cognitive Processes, 28*, 156-187.

Phillips, C. (2013b) On the nature of island constraints I: language processing and reductionist accounts. In J. Sprouse & N. Hornstein (eds.), *Experimental syntax and island effects*. Cambridge University Press.

Phillips, C., Wagers, M., & Lau, E. (2011) Grammatical illusions and selective fallibility in real-time language comprehension. In J. Runner (ed.), *Experiments at the interfaces: Syntax & Semantics, vo. 37*, pp. 153-186. Bingley, UK: Emerald.

Pickering, M. J. & Barry, G. (1993) Dependency categorial grammar and coordination. *Linguistics, 31*, 855-902.

Pickering, M. J. & Ferreira, V. S. (2008) Structural priming: a critical review. *Psychological Bulletin, 134*, 427-459.

Richards, N. (1999) Dependency formation and directionality of tree construction. *MIT Working Papers in Linguistics, 34*, 67-105.

Runner, J. T., Sussman, R. S., & Tanenhaus, M. K. (2006) Processing reflexives and pronouns in picture noun phrases. *Cognitive Science, 30*, 193-241.

Sabbagh, J. (2007) Ordering and linearizing rightward movement. *Natural Language and Linguistic Theory, 25*, 349-401.

Safir, K. (1999) Vehicle change and reconstruction in A' chains. *Linguistic Inquiry, 30*, 587-620.

Sag, I. A., Wasow, T., & Bender, E. M. (2003) *Syntactic theory: A formal introduction.* University of Chicago Press.

Schlenker, P. (1999) La flexion de l'adjectif en allemand: La morphologie de haut en bas. *Recherches Linguistiques de Vincennes, 28*.

Shan, C.-C. & Barker, C. (2006) Explaining crossover and superiority as left-to-right evaluation. *Linguistics and Philosophy, 29*, 91-134.

Shiobara, K. (2008) Prosodic phase and left-to-right structure building. *Canadian Journal of Linguistics, 53*, 253-274.

Sprouse, J., Wagers, M. W., & Phillips, C. (2012) A test of the relation between working memory capacity and syntactic island effects. *Language, 88*, 82-123.

Steedman, M. (1985) Dependency and coordination in the grammar of Dutch and English. *Language, 61*, 523-568.

Steedman, M. (1997) *Surface structure and interpretation*. Cambridge, MA: MIT Press.

Steedman, M. (2000) *The syntactic process*. Cambridge, MA: MIT Press.

Stowe, L. A. (1986) Parsing WH-constructions: Evidence for on-line gap location. *Language and Cognitive Processes, 3*, 227-245.

Stowell, T. (1981) *Origins of phrase structure.* PhD dissertation, MIT.

Sturt, P. (2003) The time course of the application of binding constraints in reference resolution. *Journal of Memory and Language, 48*, 542-562.

Sturt, P. (2007) Semantic re-interpretation and garden path recovery. *Cognition, 105*, 477-488.

Takano, Y. (1995) Predicate fronting and internal subjects. *Linguistic Inquiry, 26*, 327-340.

Torrego, E. (1984) On inversion in Spanish and some of its effects. *Linguistic Inquiry, 15*, 103-129.

Townsend, D. J. & Bever, T. G. (2001) Sentence comprehension: The integration of habits and rules. Cambridge, MA: MIT Press.

Traxler, M. J. & Pickering, M. J. (1996) Plausibility and the processing of unbounded dependencies: An eye-tracking study. *Journal of Memory and Language, 35*, 454-475.

van der Velde, F. & de Kamps, M. (2006) Neural blackboard architectures of combinatorial structures in cognition. *Behavioral and Brain Sciences, 29*, 1-72.

van Gompel, R. P. G. & Pickering, M. J. (2007) Syntactic parsing. In G. Gaskell (ed.), *The Oxford handbook of psycholinguistics*, pp. 289-307. Oxford: Oxford University Press.

Wagers, M. W., Lau, E. F., & Phillips, C. (2009) Agreement attraction in comprehension: Representations and processes. *Journal of Memory and Language, 61*, 206-237.

Wagers, M. & Phillips, C. (2009) Multiple dependencies and the role of the grammar in real-time comprehension. *Journal of Linguistics, 45*, 395-433.

Wagner, M. (2005). Prosody and recursion. PhD dissertation, MIT.

Wellwood, A., Pancheva, R., Hacquard, V., & Phillips, C. (submitted) Deconstructing a comparative illusion.

Wilder, C. (1999) Right-node raising and the LCA. In S. Bird, A. Carnie, J. D. Haugen, & P. Norquest (eds.), *Proceedings of WCCFL 18*, pp. 586-598.

Xiang, M., Dillon, B., & Phillips, C. (2009) Illusory licensing effects across dependency types: ERP evidence. *Brain and Language, 108*, 40-55.

Zwart, J.-W. (2009) Prospects for top-down derivation. *Catalan Journal of Linguistics, 8*, 161-187.