J. Linguistics **40** (2004), 571–591.
DOI: 10.1017/S0022226704002774 Printed in the United Kingdom

REVIEW ARTICLE

Foundational issues[1]

COLIN PHILLIPS & ELLEN LAU

University of Maryland

Ray Jackendoff, *Foundations of language*. Oxford: Oxford University Press, 2002. Pp. xix + 477.

1. Introduction

Ray Jackendoff is well qualified to talk about the trend towards increasing specialization and fragmentation of the science of language. For over 30 years he has been a prolific contributor to a number of different topics in syntax and semantics, and has also made notable contributions to other areas of cognitive science, including music, consciousness, spatial cognition, and psycholinguistics. It's hard to find many in the field who can match Jackendoff's breadth. Therefore, all linguists should be interested in Jackendoff's most recent book, *Foundations of language*, in which he lays out a series of objectives for the science of language, and describes some of the steps that he believes are needed in order to reintegrate theoretical linguistics with psycholinguistics and even computational neuroscience.

We find a lot to like in *Foundations of language*, and we are in broad agreement with Jackendoff on the set of goals that he proposes. In particular, we agree with his claim that the time is ripe for linguists to pay more than lip service to the long-standing mentalistic commitments of the field. However, we disagree with Jackendoff's core theoretical proposal, which is that key challenges for unifying linguistics, psycholinguistics and neuroscience are solved by his 'tripartite parallel architecture', which he describes as replacing traditionally dominant 'syntactocentric' models with independent generative systems for phonology, syntax, and semantics. We find the linguistic arguments for the architecture to be quite limited, and we are also not convinced that the architecture confers the psycholinguistic benefits that Jackendoff claims for it. We think that the real challenges for real-time

[1] We would like to thank Norbert Hornstein, Martin Pickering and an anonymous *JL* referee for valuable comments on an earlier version of this article. Preparation of this article was supported in part by grants to C.P. from the National Science Foundation (#BCS-0196004), the Human Frontier Science Program (RGY-0134), and the University of Maryland General Research Board.

deployment of linguistic knowledge are not addressed by the changes that Jackendoff proposes.

In what follows, we first lay out the overall goals of the book, and then discuss what is meant by the claim that a linguistic system is GENERATIVE, and how strong the evidence is for multiple parallel sources of generativity. We then describe what we consider to be the main challenges for bringing together linguistics, psycholinguistics and neuroscience, and what kind of architecture is needed in order to meet these challenges.

The book is divided into three sections, each comprising four chapters. Part I, 'Biological and psychological foundations', aims to explain what it means to assume a rich, innately constrained representational system for language. In this section, Jackendoff spells out in detail the kinds of combinatorial phenomena that must be captured in a theory of language, and furthermore provides a clear discussion of some of the most foundational (and most contentious) issues in current linguistic theory, such as the competence–performance distinction and the notion of Universal Grammar. Jackendoff's aim for this part of the book seems to be to make clear to the reader both the kind of problems a theory of language must explain and the reasons behind the most influential past approaches to these problems, and he succeeds in doing this in a way that is both fair-minded and entertaining to read. An important part of this section is the argument that the mentalistic commitments of 'mainstream generative linguistics' continue to be defined by the program presented at the start of Chomsky's *Aspects of the theory of syntax* (Chomsky 1965). Part II, 'Architectural foundations', is devoted to a discussion of the motivation for the parallel generative architecture, viewed from a linguistic, psycholinguistic, and evolutionary perspective. This section contains the theoretical meat of the book, where Jackendoff outlines concerns about the current state of linguistic theory and how it interfaces with other fields of cognitive science, and then goes on to present conceptual and empirical arguments for the superiority of the parallel architecture. Part III, 'Semantic and conceptual foundations', brings together a number of strands of Jackendoff's work on semantics and its position in the mental world. Although for the most part this section could be read independently from the rest of the book, as an accessible introduction to issues of semantic representation, the act of devoting a full third of the book to semantics might be seen as an implicit counter to the 'syntactocentrism' that part II critiques so heavily.

A general characteristic of the book is that the argumentation is often refreshingly even-handed. For example, the discussion in chapter 5 of how transformational and non-transformational grammars introduce similar complexity into different parts of the language system does a fine job of laying out the real issues without the rhetorical baggage that such discussions often engender. On the other hand, other parts of the book more strongly reflect Jackendoff's own biases, and it is these sections that we concentrate our attention upon for most of the remainder of this review. Given the size of the book, we cannot in any case hope to do justice to all of the issues that it touches upon. We instead focus on what we take to be the core proposal, the parallel architecture.

2. The parallel architecture

Jackendoff argues that his parallel architecture makes a crucial modification to generative grammar: it replaces the traditional 'syntactocentric' view, which localizes all generativity in the syntax, with a competence grammar that has multiple, parallel sources of generativity (125–130). Phonological, semantic, and syntactic representations are generated by separate components and connected with each other through a set of interface mapping rules, and the system is derivational only in the sense that it maps from sound to meaning or from meaning to sound when deployed in comprehension and production. Jackendoff presents this architecture in opposition to the claim that phonology and semantics are 'interpretive' systems that make do with whatever syntax provides them with. This architecture is motivated in part by the claim that it better lends itself to theories of processing and language evolution (chapters 6–8), but also by the claim that it does a better job of explaining the linguistic facts that formal theories normally concern themselves with (chapters 5 and 6). A significant portion of the book is thus devoted to a variety of interesting phenomena dealt with in Jackendoff's previous work on such topics as argument structure, idioms, and lexical semantics. Unfortunately, the connection between the rich variety of data presented and the argument for the parallel architecture is often not clear.

Our main concern with Jackendoff's linguistic arguments stems from what we perceive as ambiguity about what it means for syntax to be the only source of generativity in the grammar. Jackendoff defines SYNTACTOCENTRISM as the claim that 'all combinatoriality in language is ultimately the product of phrase structure rules' and thus that the combinatoriality of semantic and phonological representations is completely derivable from syntax (108). This claim clearly does not entail that a syntactic representation must encode every distinction present in a phonological and semantic representation. For this reason, it is hard to see the relevance of Jackendoff's extensive discussion of word-level phonological and semantic features that are not encoded in syntax (e.g. syllable structure, metrical stress, concrete vs. abstract nouns; 111–125). Similarly, the syntactocentric view does not preclude the existence of CONSTRAINTS on linguistic representations that are specific to semantics or phonology. A sentence correctly generated by the syntax may still be ill-formed because it fails to satisfy semantic or phonological constraints on well-formedness. Therefore, the existence of different types of constraints in phonology, syntax, and semantics does not decide the question of where generativity should be situated in the grammar. Also, the question of single vs. multiple sources of generativity is logically separate from the question of whether the syntax–semantic mapping is better achieved with a complex syntactic representation and simple mapping rules or simple syntax and complex mappings. Jackendoff himself emphasizes the equivalence of these two approaches in the amount of complexity introduced to the system (145f.).

Furthermore, a number of the arguments for multiple sources of generativity in language rely on Jackendoff's view that semantics is identical to conceptual structure, and that traditional subfields such as 'linguistic semantics' or 'lexical semantics' are mainly taxonomic conveniences. Since for Jackendoff all of semantics is conceptual structure, and there is no independent level of linguistic semantics, the idea that semantics is 'purely interpretive' implies that 'thought [is denied] any independent status' outside of language (108, fn. 1), a consequence that is likely to be uncomfortable for many, including Jackendoff. Of course, this argument does not hold in more standard architectures, where it is quite possible to assume that conceptual structure (a non-linguistic level) has some form of independent generative capacity, possibly even hierarchical, while still maintaining that specifically linguistic semantics is interpretive and that syntax is the sole source of the combinatoriality of language.[2]

With these preliminaries in mind, we can better evaluate the specific evidence that Jackendoff presents for multiple sources of generativity. First, Jackendoff suggests that a complex set of facts about syntax–semantics relations in argument structure 'constitute an argument against Chomsky's syntactocentric architectures', because they implicate a system that includes conceptual formation rules and complex syntax–semantics interface principles (149). The logical steps between the data and this conclusion are unfortunately not spelled out very explicitly. Jackendoff drowns the reader in interesting facts about argument structure, but provides no more than a hint of what the conceptual formation rules look like (132–149). Furthermore, he presents no evidence that the conceptual formation rules generate combinatorial LINGUISTIC representations whose structure is not isomorphic to syntactic structure. In order to differentiate his architecture from the syntactocentric alternative, Jackendoff needs to do more than point out that conceptual structure makes use of primitives that are not reducible to syntax.

To our minds the best kind of evidence for a parallel generative architecture involves cases where the phonological or semantic systems present combinatorial structures that have the same granularity as syntactic structures,

[2] Note that a recent line of research in developmental psychology presents interesting arguments that some combinatorial conceptual representations are dependent upon the prior learning of combinatorial linguistic representations (Spelke 2003, de Villiers in press). If these arguments hold up, then even the notion that conceptual structure has the properties of an 'interpretive system' becomes more plausible.

but that can be argued not to have been generated by the syntactic system. Jackendoff's discussion of Heavy NP Shift (120) provides one example that may have this character. Although the word order in (1) is unacceptable, the same sentence with a long NP object is generally considered to be quite good.

(1) *John bought yesterday a computer.
(2) John bought yesterday several pieces of hardware that he's been dreaming about for months

As Jackendoff notes, this acceptability contrast has previously been argued to be rooted in processing load constraints on production or comprehension (Hawkins 1994; Gibson 1998; Wasow 2002). Another (not mutually incompatible) explanation for this contrast in terms of the parallel architecture is that the word order in (2) displays the prosodic structure generated by the phonological component on the basis of constraints like 'put the longest intonational phrase at the end', while rules of the syntactic component such as 'don't put temporal adverbs before the direct object' are overridden. This seems to be what Jackendoff has in mind when he says this is a situation where 'the needs of prosody are forcing a non-optimal syntactic structure' (121). Another similar argument could be constructed based upon second-clitic phenomena in Serbo-Croatian, as discussed in some of Jackendoff's earlier work (Jackendoff 1997). However, Jackendoff's formulation of the Heavy NP Shift (HNPS) phenomenon remains extremely close to a standard syntactocentric approach, in which both the default word order and the shifted order are generated by the syntactic component, and the choice is regulated by phonological constraints (e.g. Zec & Inkelas 1990). If the needs of prosody force a 'non-optimal syntactic structure' (121) under Jackendoff's account, then it seems that the syntactic component must still generate both word orders. Unfortunately, Jackendoff does not provide additional information on how his account differs from a non-transformational syntactocentric approach to HNPS, and neither does he provide any empirical evidence that would favour his approach. More generally, it is interesting to note that the cases where the word order favoured by phonology diverges from that favoured by syntax appear to be quite rare. So rare, in fact, that Jackendoff postulates a correspondence rule that requires morphophonology to preserve the order of syntactic constituents (118), although he notes that this may just be a strong default rather than an absolute rule.

A further type of evidence for multiple generative components involves situations where the representations required by other components cannot be mapped from the syntactic structure in any obvious way. Jackendoff suggests that the well-known mismatch between intonational phrases and syntactic structures provides this kind of argument for his architecture (118–121). If the intonational phrasing of a sentence exhibits a structural bracketing that cannot be simply related to the syntactic structure of the sentence, the phonological component must have some level of generativity.

But while Jackendoff is correct in his claim that this kind of mismatch is entirely consistent with the parallel architecture, there are also ways to map straightforwardly from syntactic structure to intonational structure, particularly in syntactic approaches that allow incremental left-to-right syntactic composition (e.g. Phillips 1996, 2003; Steedman 2000b). For example, the classic example in (3) shows a clear mismatch between syntactic and intonational structure when the structures are presented as a whole, but once we consider how the sentence could be assembled incrementally, we find that whenever a new intonational phrase is formed it corresponds to a current syntactic constituent (4). A much fuller discussion along related lines can be found in Steedman (2000b).

(3) (a) *Syntax*
[$_{NP}$ this] [$_{VP}$ is [$_{NP}$ the cat [$_{CP}$ that [$_{VP}$ caught [$_{NP}$ the rat
[$_{CP}$ that [$_{VP}$ stole [$_{NP}$ the cheese]]]]]]]]

(b) *Phonology*
[$_{IntP}$ this is the cat] [$_{IntP}$ that caught the rat] [$_{IntP}$ that stole the cheese]

(4) (a) *Syntax*
[$_{CP}$ [$_{NP}$ this] [$_{VP}$ is [$_{NP}$ the cat]]]
Phonology
[$_{IntP}$ this is the cat]

(b) *Syntax*
[$_{CP}$ [$_{NP}$ this] [$_{VP}$ is [$_{NP}$ the cat [$_{CP}$ that [$_{VP}$ caught [$_{NP}$ the rat]]]]]]
Phonology
[$_{IntP}$ this is the cat] [$_{IntP}$ that caught the rat]

Therefore, the syntax–phonology 'mismatches' could implicate a need for an entirely independent system for bracketing sentences in the phonology, or they could just implicate the need to rethink the dynamics of how syntactic structures are mapped onto phonological structures. Jackendoff is right that the phonological phrasing in examples like (3b) does not immediately follow from the bottom-up derivations of approaches like the Minimalist Program. However, even in well-known approaches to intonational phrasing that accord the phonology a good deal of independence, intonational phrases are mostly derived by projecting bracketings from syntactic structures, and then applying constraints on phrase length and 'balancing' to the output of these algorithms (e.g. Selkirk 1984; Nespor & Vogel 1986), and thus are still largely syntactocentric.

In principle, a fourth type of evidence against syntax as the source of generativity in language might come from a demonstration of phonology–semantics relations that are not represented at all in the syntax. Along these lines, Jackendoff suggests that his architecture allows a natural account of the relation between semantics and phonology in focus phenomena, in contrast to traditional syntactic frameworks, which must stipulate some invisible representation of focus in the syntax (409f.). However, the direct phonology–semantics interface does little work beyond focus phenomena (as Jackendoff concedes, 203, fn. 3), and in particular does not seem to support syntax-independent COMBINATORIAL operations (131f.). Thus, the introduction of a direct phonology–semantics interface in order to account for focus phenomena adds just as much complexity to the system as the use of a syntactic focus feature in previous treatments. The issue must be decided by empirical arguments, then, rather than by parsimony considerations.

The phenomena that Jackendoff presents as evidence for his parallel architecture (argument structure, Heavy NP Shift, intonational phrasing, focus) are all undeniably interesting, but it is disappointing to find that these are the only cases in the book that seem to even have the potential to differentiate the architecture from its syntactocentric alternatives, and in each case Jackendoff provides little in terms of arguments for why his account should be favoured over alternative approaches. This leaves a somewhat slender empirical base for a new theory to rest on. To his credit, Jackendoff frequently admits that more work is needed to support the architecture and offers it as a challenge for the future. However, since the parallel architecture is presented throughout the book as offering the field a revolutionary perspective, some more solid motivation for this major theoretical shift would be welcome. *Foundations of language* does a good job of illustrating to readers from other fields the complexity of the phenomena that a serious account of language must deal with. Unfortunately, it does a rather poorer job of showcasing the careful empirical argumentation that characterizes much of the best work in linguistics, including some of Jackendoff's own work.

We should point out that we are not particularly committed to the syntactocentric architecture. We agree with Jackendoff on the importance of the question of whether there are multiple sources of generativity (at the same granularity) in language. However, barring more compelling arguments, we see little reason for advocates of syntactocentric approaches to abandon their architecture, based on the linguistic evidence that Jackendoff presents.

3. Challenges for unification

3.1 *Real-time systems*

In addition to the question of increased empirical coverage, an important part of Jackendoff's argument is the claim that his parallel architecture has significant advantages for the goal of uniting linguistics, psycholinguistics and neurolinguistics. This argument for the architecture is, of course, only relevant to the extent that the goal of better integrating these subfields is well motivated; but on this basic point, at least, we are in complete agreement with Jackendoff. However, it is important to review the structure of Jackendoff's argument about the psycholinguistic import of the architecture, since the discussion in the book leaves out some important steps.

Jackendoff argues that the parallel architecture 'offers a theoretical perspective that unifies linguistics with psycholinguistics more satisfactorily than has previously been possible' (196). This claim is based on a criticism of the standard architecture that 'begins with phrase structure composition and lexical insertion, and branches outward to phonology and semantics', something that is 'quite at odds with the logical directionality of processing, where speech perception has to get from sounds to meaning, and speech production has to get from meanings to sounds' (197). In contrast with this, 'the parallel constraint-based architecture is logically NON-directional: one can start with any piece of structure in any component and pass along logical pathways provided by the constraints to construct a coherent larger structure around it … Because the grammar is logically non-directional, it is not inherently biased toward either perception or production – unlike the syntactocentric architecture, which is inherently biased against both' (198). This makes it possible to 'describe the logic of processing in terms isomorphic to the rule types in the parallel grammar' (199).

This is the extent of Jackendoff's psycholinguistic argument. He provides no more specific examples of where the parallel architecture succeeds and an architecture based on transformational grammar fails. The line of reasoning sounds attractive at first, but we think that it becomes rather less compelling once it is unpacked in more detail. The argument seems to imply the claims in (5).

(5) (a) The structures that speakers build in real time faithfully reflect their grammatical knowledge.
 (b) Standard architectures based on transformational grammar do not allow real-time building of grammatical structures.
 (c) The tripartite non-derivational architecture does make it possible to build grammatical structures in real time.
 (d) Parsing and production are mirror images of one another, they are unidirectional mappings between sound and meaning, and they do not rely on syntax to regulate structure building in phonology and semantics.

The first step of this argument takes a controversial position, though we are inclined to agree with Jackendoff on this point. On the other hand, we are in less agreement with the remainder of the argument.

(a) *Are real-time language processes grammatically accurate?* If the generalizations described in linguists' grammars accurately capture the range of representations that speakers are able to construct when freed from time constraints, but do not do a good job of capturing the range of structures that speakers construct in real time, then there is little point in trying to construct a theory that predicts that grammatical knowledge is available in real time. This issue has been addressed sporadically since the 1960s. In two classic surveys of psycholinguistics published 30 years ago, Fodor, Bever & Garrett (1974) and Levelt (1974) conclude that what is constructed in real time is a good approximation to the surface syntactic structures of the time, but are skeptical of the notion that anything like the operations of a transformational grammar are used to transform these structures into deep structures. Following the massive expansion of linguistics and psycholinguistics since that time, many on both sides have come to assume that the structures built in the first parse provide only a rough-and-ready approximation to what the grammar allows, and that much of the rich detail of grammar is only available more slowly, as the result of a separate structure-building process. This conclusion is based on considerations such as the existence of garden-path sentences and other types of processing breakdown, the slowness and subtlety of many grammaticality judgments, and on the existence of many sentences that are readily parsed and understood but judged to be ungrammatical (e.g. Townsend & Bever 2001; Ferreira 2003).[3] Despite this, we are inclined to believe that Jackendoff is correct to assume that grammatical knowledge is deployed very rapidly in language processing. There is now a good deal of experimental evidence for faithfulness to grammatical constraints in real-time processing, in areas such as argument structure (e.g. Boland, Tanenhaus, Garnsey & Carlson 1995; Pickering & Traxler 2001), constraints on long-distance dependencies (Stowe 1986; Traxler & Pickering 1996; McElree & Griffith 1998; Phillips 2004b), and binding constraints (e.g. Nicol & Swinney 1989; Sturt 2003; Kazanina, Lau, Lieberman, Phillips & Yoshida 2004). Further support comes from the consistent finding in ERP studies on sentence processing that syntactic violations of many different kinds elicit specific electrophysiological responses within 300–600 ms of the appearance of the offending word (for recent reviews see Hagoort, Brown & Osterhout 1999; Friederici 2001). Meanwhile, garden-path sentences and the difficulty of multiply centre-embedded structures do not challenge the AVAILABILITY of grammatical knowledge in real time, but merely show that use of this knowledge is constrained by memory resource limitations and the ambiguity of the input. Finally, the phenomenon of parsable-but-ungrammatical sentences does not imply a parser that ignores the rules of the grammar. On the contrary, when people encounter such sentences they are generally able to pinpoint the source of the deviance very easily, suggesting that they are not ignoring the rules of the grammar. For further discussion of all of these issues see Phillips (2004a).

[3] In the literature on language production, it has remained more common to assume that the structures built in real time are grammatically faithful (cf. Bock 1995).

On balance, then, we think that real-time language processes show great grammatical precision, and that linguists and psycholinguists alike should be held accountable for the successes of real-time linguistic computation just as they are held accountable for the successes of language acquisition. Therefore we consider Jackendoff's assumption (a) to be quite justified.

(b) *Is transformational grammar compatible with parsing and production?* On the other hand, the second step in the argument, the claim that transformational grammars are fundamentally incompatible with real-time structure-building, is a well-known criticism that once was valid but that has been overtaken by developments in the grammatical theory. In an *Aspects*-style transformational grammar it was only possible to assemble a surface structure by first generating a deep structure and then optionally applying transformations to generate the surface structure. These transformations were not necessarily structure-preserving. Jackendoff's argument works for a grammar of this type, and it is identical to an argument developed in detail by Fodor et al. (1974), who showed that the surface structure configurations of contemporary grammars were only justified with reference to the deep structure-rules and the transformations that generated them. However, the situation is rather different in recent transformational grammars that share the assumption that all structural configurations at all levels are motivated by specific grammatical requirements (e.g. Chomsky 1995). This approach shares with non-transformational theories such as HPSG and Categorial Grammar the property that one can directly evaluate the well-formedness of surface configurations (Pollard & Sag 1995; Steedman 2000a). Since all local surface structure configurations are motivated by specific feature-checking requirements, it should be possible to incrementally evaluate the well-formedness of individual pieces of structure as a sentence is built up. This makes it possible to turn the derivations around and map from surface structures to underlying structures with no negative impact on the coverage of the grammar. There are some arguments that it may even be preferable to turn around the derivations in this way (e.g. Phillips 1996, 2003; Richards 1999, 2002), but it is certainly possible in principle to run current transformational grammar derivations in either direction,[4] thereby neutralizing the force of Jackendoff's argument.

(c) *Does the parallel architecture facilitate incremental structure building?* The third step of Jackendoff's argument involves the assumption that a

[4] Note that in order for this claim about feature-based transformational grammars to succeed, it must be the case that the operations known as 'PF movement' either do not exist or are both structure-preserving and motivated by specific featural requirements. Although existing theories are sometimes unclear on this issue, we see no principled barrier to this approach. Note also that satisfaction of featural requirements in surface structure does not, of course, guarantee well-formedness in a transformational grammar, since it is also necessary to map surface structures onto other levels of representation, including some version of 'underlying' structure.

non-directional grammar straightforwardly enables the assembly of grammatically faithful structures based on partial input. We agree that the ability to build structures incrementally and accurately is one of the primary challenges for combining linguistic theories with psycholinguistic models, but it is not clear to us that Jackendoff's architecture has any advantage over others in dealing with this challenge. Since we think that the problem of incremental parsing is a crucial one, however, we examine the relevant issues in more detail.

The most glaring obstacle to capturing incremental structure building with current grammars is that grammatical theories typically provide accounts of which combinations of words form well-formed constituents, but the sequences of words that must be combined in real time often do not correspond to complete constituents. If real-time processes build only well-formed structures (see assumption (5a)), and if the grammar defines well-formedness in terms of possible constituents, then we might expect that structure-building should occur only when a string is identified that forms a possible constituent. However, in a right-branching sentence structure like (6), no constituents are completed until the last word of the sentence, yet speakers clearly do not need to wait until the end of the sentence to begin compositional interpretation. This basic fact does not follow immediately from a standard phrase structure grammar, directional or otherwise.

(6) [John [will [say [that [Mary [hopes [that [Sally [likes [the pizza]]]]]]]]]

One approach that can handle the problem posed by (6) modifies the grammatical theory of syntactic composition to allow 'non-standard' constituents, such that strings like 'John will say' are treated as well-formed constituents and incremental assembly of (6) becomes straightforward. This approach can be justified by showing that these sequences also behave like constituents according to other, more traditional diagnostics (Steedman 2000a; Phillips 2003). However, this move only addresses one part of the problem, since there are strings of words that appear to allow incremental interpretation but are unlikely to qualify as constituents under even the wildest grammatical theories. In particular, the results of some of our own experiments have forced us to reconsider the extent to which non-standard constituency will solve the incrementality problem. We will present one example of this in some detail, as it illustrates the point that an account of accurate incremental interpretation seems to require predictive structure building, and that it is not feasible to simply rely on the identification of complete well-formed constituents in the input.

Our example comes from recent experiments on parsing in Japanese, a language whose strongly head-final property presents notorious challenges for incremental structure building. Japanese relative clauses are head-final, creating the possibility for various types of temporary syntactic ambiguities (Inoue 1991; Mazuka & Itoh 1995), as in (7). Here a comprehender would have no indication that s/he is dealing with a relative clause until the head of the relative clause is encountered, leading to processing difficulty.

(7)	John-ga	Mary-ni	[$_{RC}$ __$_i$	ringo-o	tabeta]	inu$_i$-o	ageta.
	John-NOM	Mary-DAT		apple-ACC	eat.PST	dog-ACC	gave

'John gave Mary a dog that ate an apple.'

However, by inserting a numeral classifier that mismatches the following noun, Japanese speakers are able to recognize the presence of an upcoming relative clause. The grammar of Japanese allows a genitive numeral classifier associated with the head of a relative clause to precede the relative clause, and it follows from this that Japanese allows grammatical strings in which a numeral classifier and the following noun are mismatched, since the noun may be the subject of the relative clause. The fact that the grammar allows this does not automatically entail that Japanese speakers should be able to immediately detect a relative clause structure when they encounter the classifier–noun mismatch, yet in fact this is exactly what they seem to do (Yoshida, Aoshima & Phillips 2004). When speakers were asked to generate written completions for sentence fragments with classifier–noun matches, (8a), only 0.1% of the completions contained relative clauses, but when the fragments involved classifier–noun mismatches, (8b), the proportion of relative clause completions increased to 81%.

(8)	(a)	Dono NP-ni	NP-TOP san-*nin*-no	Adj	*sensee*-ga …
		which NP-DAT	**THREE-CLASSIFIER(human)-GEN**		teacher-NOM
	(b)	Dono NP-ni	san-*satu*-no	Adj	*sensee*-ga …
		which NP-DAT	**THREE-CLASSIFIER(book)-GEN**		teacher-NOM

The numeral classifier cue was similarly effective in an on-line self-paced reading experiment. This study compared word-by-word reading times in sentences that contained a relative clause, but with either a matching or mismatching classifier–noun sequence at the start of the relative clause. In the version with a classifier–noun mismatch, reading times at the end of the relative clause showed facilitation relative to the classifier–noun match conditions, suggesting that speakers had already figured out that they were parsing a relative clause structure.

Finally, in order to test whether Japanese speakers were able to compute the consequences of building a relative clause structure, an additional self-paced reading study tested whether the classifier–noun mismatches would be sufficient to introduce effects of island constraints on long-distance scrambling. Earlier studies on Japanese parsing showed that speakers favour a long-distance scrambling analysis of fronted dative NPs (Aoshima, Phillips & Weinberg 2004). This preference is potentially in conflict with the constraint on scrambling out of relative clauses (Saito 1985). This study replicated Aoshima et al.'s finding of a pre-verbal 'Filled Gap Effect' in sentences with matching numeral classifiers, (9a, b), but showed that this effect disappeared in conditions where mismatching classifiers indicate an upcoming relative clause structure, (9c, d). The Filled Gap Effect (Crain & Fodor 1985; Stowe 1986) was observed in a slowdown immediately after the embedded dative NP in the scrambled-match condition (9a) relative to the unscrambled-match condition (9b). This effect is expected if the fronted dative NP in (9a) undergoes long-distance scrambling that places it inside the embedded clause whereas the unscrambled dative NP in (9b) is interpreted in-situ (cf. Kamide & Mitchell 1999). No such contrast is found in the classifier-mismatch conditions (9c, d), indicating that long-distance scrambling is blocked when classifiers cue an upcoming relative clause.

(9) (a) *Scrambled-Match*
Dono-seeto-ni tannin-wa [[3-nin-no tosioita sensee]-ga
which-student-DAT class-teacher 3-CL-GEN aged teacher-NOM
atarasii koochoo-ni yorokonde okutta] hon-o
new president-DAT gladly gave book-ACC
yomasemasita-ka
read-made-HONORIFIC-**Q**

(b) *Unscrambled-Match*
Tannin-wa dono-seeto-ni [[3-nin-no tosioita sensee]-ga
class-teacher which-student-DAT 3-CL-GEN aged teacher-NOM
atarasii koochoo-ni yorokonde okutta] hon-o
new president-DAT gladly gave book-ACC
yomasemasita-ka
read-made-HONORIFIC-**Q**
'Which student did the class teacher force to read a book that three old teachers gladly gave to the new president?'

(c) *Scrambled-Mismatch*
Dono-seeto-ni tannin-wa 3-satu-no [[tosioita sensee-ga]
which-student-DAT class-teacher 3-CL-GEN aged teacher-NOM
atarasii koochoo-ni yorokonde okutta] hon-o
new president-DAT gladly gave book-ACC
yomasemasita-ka?
read-made-HONORIFIC-**Q**

(d) *Unscrambled-Mismatch*
Tannin-wa dono-seeto-ni 3-satu-no [[tosioita sensee-ga]
class-teacher-TOP which-student-DAT 3-CL-GEN aged teacher-NOM
atarasii koochoo-ni yorokonde okutta] hon-o
new president-DAT gladly gave book-ACC
yomasemasita-ka?
read-made-HONORIFIC-**Q**
'Which student did the class teacher force to read three books that the old teacher gladly gave to the new president?'

The implication of this is that speakers are able to build a full sentence skeleton containing a main clause and a relative clause, based on having encountered just a couple of subject NPs and a mismatching numeral classifier. Sequences such as this do not form a syntactic constituent under any approach that we are aware of, and thus it seems that this structure must be built with a predictive mechanism that uses already-seen material to project not-yet-encountered material. This property does not come for free even in grammars that allow incremental left-to-right formation of constituents. Findings of this kind are interesting in the respect that they show that speakers are able to build structure incrementally and accurately, even when presented with sequences of words that cannot be combined directly using the rules of the grammar, but they also beg the question of HOW grammatical accuracy is preserved once predictive mechanisms are deployed. Predictive structure building is a hallmark of Left Corner Parsers (e.g. Griffiths & Petrick 1965; Abney & Johnson 1991), but a long-standing challenge for this type of parser has been to constrain the predictive mechanisms in such a way that they do not overgenerate possible structures.

What the above discussion should make clear is that, while it may be possible to incorporate the components of Jackendoff's architecture into a real-time system that is able to carry out constrained predictive structure-building, the ability to do this is not something that comes for free in Jackendoff's system any more than it does in most other current theoretical frameworks. Schneider (1999) and Aoshima (2003) argue that lexicalized grammar formalisms such as Minimalism, HPSG, and Categorial Grammar offer advantages for accurate predictive structure-building over grammars based on phrase structure rules and/or construction templates, but the structure building mechanisms still do not follow automatically from the grammar formalism. A rare example of a theoretical approach that does seem to straightforwardly imply accurate predictive structure building can be found in the 'Dynamic Syntax' approach of Kempson and colleagues (Kempson, Meyer-Viol & Gabbay 2001).

(d) *Information flow among levels*. Jackendoff argues that his architecture is more suitable for parsing and production because of its lack of commitment to a direction of information flow among levels of representation. He makes the provisional assumption that parsing and production are the mirror image of one another, mapping from phonology to syntax to semantics (parsing) or from semantics to syntax to phonology (production), and argues that a 'syntactocentric' architecture is compatible with neither of these (196–198).

Of course, few people would argue that in comprehension or production, a speaker actually starts at the syntactic level. It is certainly true that comprehension requires a mapping from sound to meaning and that production requires a mapping from meaning to sound, but this fact places few *a priori* constraints on the flow of information among the specifically linguistic levels of representation in comprehension and production, i.e. phonology, syntax, and semantics. This is because neither of the endpoints of parsing and production is specifically linguistic in nature, and because it is clear that representations at all levels are constructed incrementally during parsing and production. Therefore, in order to show, for example, that phonology is independently generative as implied by Jackendoff's architecture, more sophisticated arguments are needed.

In the case of comprehension, the crucial question is whether the bits of phonology being passed to syntax and semantics get some kind of independent organizing structure from phonology or semantics (above the word level) before reaching the syntactic component. If the pieces of phonology being sent to the syntactic and semantic components are word-sized (or smaller), then the direction of information flow would still be consistent with syntactic theories that assume semantics and phonology are interpretive with respect to the structure provided by the syntactic component. It would be hard to prove experimentally that the size of the phonological chunks being sent to syntax couldn't EVER be large. What can and has been shown is that in at least some cases, phonological information is informing syntactic analysis early enough that there is time for the syntactic information to feed back into the perception of the word before the word is completed. Marslen-Wilson (1975) showed that the degree of syntactic and semantic contextual fit of the first syllable of a word impacted the number of corrections to errors in the second syllable of the word. Thus we know that it is possible for smaller-than-word-size pieces of phonology to be processed by the syntactic component, consistent with either the syntactocentric account or Jackendoff's account. One would need to show that it is possible for multi-word units to be structured phonologically before any syntactic combinatorial information is accessed in order to use the directionality of processing to motivate Jackendoff's parallel architecture.

In the case of language production, the logic of the issue is similar, though with a couple of additional twists. The crucial question is whether combinatorial semantic representations are built independent of syntax, and whether such representations interface directly with syntactic structures. If the pieces of semantics that are fed to syntax and phonology in production are word-sized (or smaller), then the direction of information flow among levels would still be compatible with the syntactocentric architecture. It is widely assumed that conceptual structure (the 'language of thought') consists of compositional representations that have the granularity of propositions, and that these exist independent of language (but see Spelke 2003 for an alternative view). It is also generally assumed that language production involves the mapping of compositional conceptual structures onto compositional linguistic representations. These standard assumptions about language production are not themselves decisive regarding the directionality of information flow within the language system, since they make no claim about the relation between syntax and LINGUISTIC semantic representations. As discussed earlier, Jackendoff denies the distinction between conceptual structure and a specifically linguistic semantic level (chapter 10), and so effectively sidesteps the question of whether linguistic semantics is an interpretive system or an independent generative system; his assumption that conceptual structure is an independently generative system is relatively uncontroversial. If Jackendoff is correct in arguing for a single system for representing meaning, changes will clearly need to be made to the many syntactocentric architectures which assume such a distinct (interpretive) linguistic semantics. However, to use facts about information flow from production as a separate and independent argument against these syntactocentric architectures, one would have to show that well-formed representations at this linguistic semantic level (and not the conceptual level) precede well-formed syntactic representations during production. We are not aware of work in the production literature that would speak to this question.

In sum, while we agree with Jackendoff's objectives for combining the concerns of theoretical linguistics and psycholinguistics, we consider the 'non-directionality' of the parallel architecture to be a questionable virtue, and would argue instead that the main challenge for unification in this area involves the question of how to build structures accurately and incrementally in real time. This challenge could be viewed as the 'Logical Problem of Language Processing', and it remains somewhat mysterious under most theoretical approaches.

3.2 *Neural encoding*

With regard to the challenge of grounding linguistics in neuroscience, Jackendoff lays out an interesting set of 'Four Challenges for Cognitive Neuroscience' (section 3.5). He recognizes that the main challenge is to explain how compositional structure is encoded in groups of neurons, rather than one of brain localization, and he presents the challenges in ecumenical terms that are relevant to most linguistic theories. Two of the four challenges will be familiar to most linguists: 'the massiveness of the binding problem' refers to the need to recognize the richness of linguistic combinatorial systems; 'the problem of variables' refers to rule-based behaviours. The two other challenges involve the distinction between long-term and short-term representations of language structure, a topic that linguists worry about less often. The 'problem of 2' refers to the problem of how to maintain multiple representations of the same lexical item, as in 'the black cat scared the white cat', in a system in which there is a single long-term neural representation of each word. The solution here may require the ability to perform the neural equivalent of copying the contents of an item in long-term memory into multiple registers in a computational workspace. This is far from trivial, since it remains unclear how the equivalent of a computer register could be reconstructed in a neural architecture.

The most interesting of the four challenges for cognitive neuroscience is the problem of 'binding in working memory vs. long-term memory'. It is clear that speakers must store long-term representations of words, and it is generally assumed that these representations must be encoded through patterns of synaptic connectivity among neurons, which are easily maintained over long periods of time (cf. Pulvermüller 2003). It is also clear that speakers must be able to create short-term representations of structured combinations of words, i.e. phrases and sentences, and that these representations must be constructed on a time-scale that is too fast for encoding structure via changes in synaptic connection strengths, i.e. hundreds of milliseconds. This has spawned an interesting and growing literature on mechanisms for building and encoding structure in neurons on a short time-scale (e.g. Shastri & Ajjanagadde 1993; Hummel & Holyoak 1997; Whitney & Weinberg 2003). Jackendoff emphasizes, however, that such mechanisms are unlikely to be sufficient. Although most sentences are not memorized and must be constructed on the fly, speakers are also able to store larger phrases, sentences, and even larger discourses in long-term memory, and this calls for mechanisms that encode structures using patterns of synaptic connectivity. This in turn suggests a need for two different ways of neurally encoding any linguistic structure, and for a way to translate between the two representational schemes. If true, this is a very important point, since it implies that there is no single answer to the question 'how are sentences encoded in the brain', and further suggests that learners must – somewhat redundantly – master two representational systems for their target language. The one way of avoiding this conclusion is to argue that long-term memorized sentences are simply encoded as strings of words, and that sentences are organized into hierarchical structures only when they are transferred to short-term encoding or 'working memory'. To our knowledge, there is little current evidence on the form of long-term representations of sentences, and this is an important question for future research to address. Note that Jackendoff's arguments about memorized phrases and sentences also play an important role in his treatment of the content of the lexicon. As argued in detail in chapter 6 ('Lexical storage vs. on-line computation'), the lexicon contains larger structural units than is assumed in many current theories.

We would add to Jackendoff's list of challenges a fifth challenge for cognitive neuroscience, which we call the 'problem of discreteness'. In chapter 2 Jackendoff acknowledges the concern that the currency of linguistic models is typically discrete categories, whereas most well-known phenomena at the neurophysiological level appear to be continuously varying, non-discrete phenomena. (A single spike of a neuron is admittedly a discrete event, but it is generally assumed that information is encoded in the temporal patterns of spikes, rather than in individual spikes, Riecke, Warland, deRuyter van

Steveninck & Bialek 1996.) Jackendoff suggests that this discrepancy between linguistics and neuroscience be resolved by giving up the notion that linguistic categories are genuinely discrete. We suspect that this move is premature and based on conflation of different notions of 'category'. For example, Jackendoff points to the well-known finding that 'categorical perception of phonemes is not absolutely sharp' (25), with a narrow range of uncertainty and context sensitivity around the acoustic boundary between categories such as /b/ and /p/. A similar argument involving semantic categories for colour is presented in chapter 11. However, it is a mistake to infer from uncertainty in DECISIONS about category membership that category REPRESENTATIONS are themselves non-discrete. It is quite common in everyday life to encounter uncertainty in decisions where we know that we face a binary choice, and thus the uncertainty seen in phonetic identification tasks does not provide evidence that the categories are non-discrete. The same is true in phonology. Research on speech perception has shown that many different factors can affect decisions about what phonological category to identify with particular acoustic stimuli (Diehl, Lotto & Holt 2004), but the phonological categories themselves behave like discrete computational units. Phonological processes (e.g. epenthesis, syllabification, stress assignment) apply to all instances of an appropriate category, with no regard to the 'goodness' of the instance of the category. Storage of words in long-term memory requires only an encoding of the phonological categories of the word. Even if it is possible to store additional phonetic detail with the lexical representations of certain words, this is by no means necessary, as shown by the fact that novel words can be learned from reading text. We therefore suggest that the encoding of discrete category representations be added as a fifth challenge for cognitive neuroscience.

We should point out that although Jackendoff presents the list of challenges as if they have been largely ignored by the computational neuroscience community, the situation is actually more encouraging if one looks beyond the eliminative connectionist tradition that has attracted the most attention from linguists (e.g. Elman, Bates, Johnson, Karmiloff-Smith, Parisi & Plunkett 1996). There have been a growing number of attempts to design systems that can parse and represent complex linguistic structures while respecting neural constraints (e.g. Pulvermüller 2003; Whitney & Weinberg 2003). One mark of the value of Jackendoff's presentation of the challenges is a recent neuro-computational proposal that is explicitly organized as a set of answers to Jackendoff's challenges (van der Velde & de Kamps 2004).

4. Conclusion

We remain unconvinced that the parallel architecture solves all of the problems that Jackendoff claims that it solves. Nevertheless, we should emphasize that we are in broad agreement with Jackendoff that the issues addressed by *Foundations of language* are the right kinds of issues to be concerned with. We would be delighted if books of this kind were to encourage more linguists to engage in discussions of foundational issues with psychologists, philosophers, and neuroscientists.

Jackendoff makes an interesting claim in chapter 2 about the field's continued reliance on the program laid out in chapter 1 of *Aspects of the theory of syntax* (Chomsky 1965). He argues that what counted as a reasonably explicit set of mentalistic commitments in 1965 no longer cuts it, in light of advances in psycholinguistics and neuroscience in the intervening 40 years. Most linguists would agree that their theories aim to characterize mental phenomena, but it is surprisingly difficult to find more explicit discussion of what competence grammars are supposed to be theories of. Jackendoff repeatedly claims that they are intended as 'metaphors', though many linguists are likely to take umbrage at such a characterization, and would want their theories to make stronger claims than that. If *Foundations of language* stimulates linguists to try to do a better job of answering the challenges that Jackendoff addresses, then it will have done a truly useful service to the field.

REFERENCES

Abney, S. & Johnson, M. (1991). Memory requirements and local ambiguities of parsing strategies. *Journal of Psycholinguistic Research* **20**. 233–250.

Aoshima, S. (2003). *The parsing and grammar of* wh-*dependencies*. Ph.D. dissertation, University of Maryland.

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

Bock, J. K. (1995). Sentence production: from mind to mouth. In Miller, J. L. & Eimas, P. D. (eds.), *Speech, language, and communication*. San Diego: Academic Press. 181–216.

Boland, J., Tanenhaus, M. K., Garnsey, S. M. & Carlson, G. N. (1995). Verb argument structure in parsing and interpretation: evidence from *wh*-questions. *Journal of Memory and Language* **34**. 774–806.

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

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

Crain, S. & Fodor, J. D. (1985). How can grammars help parsers? In Dowty, D., Karttunen, L. & Zwicky, A. M. (eds.), *Natural language parsing: psycholinguistic, computational, and theoretical perspectives*. Cambridge: Cambridge University Press. 94–128.

de Villiers, J. G. (in press). Can language acquisition give children a point of view? In Astington, J. W. & Baird, J. A. (eds.), *Why language matters for theory of mind*. Oxford: Oxford University Press.

Diehl, R. L., Lotto, A. J. & Holt, L. L. (2004). Speech perception. *Annual Review of Psychology* **55**. 149–179.

Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D. & Plunkett, K. (1996). *Rethinking innateness: a connectionist perspective on development*. Cambridge, MA: MIT Press.

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

Fodor, J. A., Bever, T. G. & Garrett, M. F. (1974). *The psychology of language*. New York: Holt, Rinehart & Winston.

Friederici, A. (2001). Towards a neural basis of auditory sentence processing. *Trends in Cognitive Sciences* **6**. 78–84.

Gibson, E. (1998). Linguistic complexity: locality of syntactic dependencies. *Cognition* **68**. 1–76.
Griffiths, T. V. & Petrick, S. R. (1965). On the relative efficiencies of context free grammar recognizers. *Communications of the Association for Computing Machinery* **8**. 289–300.
Hagoort, P., Brown, C. & Osterhout, L. (1999). The neurocognition of syntactic processing. In Brown, C. M. & Hagoort, P. (eds.), *Neurocognition of language*. Oxford: Oxford University Press. 273–316.
Hawkins, J. A. (1994). *A performance theory of order and constituency*. Cambridge: Cambridge University Press.
Hummel, J. & Holyoak, K. (1997). Distributed representations of structure: a theory of analogical access and mapping. *Psychological Review* **104**. 427–467.
Inoue, A. (1991). *A comparative study of parsing in Japanese and English*. Ph.D. dissertation, University of Connecticut.
Jackendoff, R. (1997). *The architecture of the language faculty*. Cambridge, MA: MIT Press.
Kamide, Y. & Mitchell, D. C. (1999). Incremental pre-head attachment in Japanese parsing. *Language and Cognitive Processes* **14**. 631–662.
Kazanina, N., Lau, E., Lieberman, M., Phillips, C. & Yoshida, M. (2004). Use of grammatical constraints in the processing of backwards anaphora. Talk at the 17th Annual CUNY Sentence Processing Conference, College Park, MD.
Kempson, R., Meyer-Viol, W. & Gabbay, D. (2001). *Dynamic Syntax*. Oxford: Blackwell.
Levelt, W. J. M. (1974). *Formal grammars in linguistics and psycholinguistics* (3 vols.). The Hague: Mouton.
Marslen-Wilson, W. D. (1975). Sentence perception as an interactive parallel process. *Science* **189**. 226–228.
Mazuka, R. & Itoh, K. (1995). Can Japanese speakers be led down the garden path? In Mazuka, R. & Nagai, N. (eds.), *Japanese sentence processing*. Hillsdale, NJ: Lawrence Erlbaum Associates. 295–329.
McElree, B. & Griffith, T. (1998). Structural and lexical constraints on filling gaps during sentence processing: a time-course analysis. *Journal of Experimental Psychology: Learning, Memory, & Cognition* **24**. 432–460.
Nespor, M. & Vogel, I. (1986). *Prosodic phonology*. Dordrecht: Reidel.
Nicol, J. L. & Swinney, D. (1989). The role of structure in coreference assignment during sentence comprehension. *Journal of Psycholinguistic Research* **18**. 5–20.
Phillips, C. (1996). *Order and structure*. Ph.D. dissertation, MIT.
Phillips, C. (2003). Linear order and constituency. *Linguistic Inquiry* **34**. 37–90.
Phillips, C. (2004a). Linguistics and linking problems. In Rice, M. & Warren, S. (eds.), *Developmental language disorders: from phenotypes to etiologies*. Mahwah, NJ: Erlbaum. 241–287.
Phillips, C. (2004b). The real-time status of island constraints. Ms., University of Maryland. [Submitted.]
Pickering, M. J. & Traxler, M. J. (2001). Strategies for processing unbounded dependencies: lexical information and verb-argument assignment. *Journal of Experimental Psychology: Learning, Memory, & Cognition* **27**. 1401–1410.
Pollard, C. & Sag, I. A. (1995). *Head-driven phrase structure grammar*. Chicago: University of Chicago Press.
Pulvermüller, F. (2003). *The neuroscience of language: on brain circuits of words and serial order*. Cambridge: Cambridge University Press.
Richards, N. (1999). Dependency formation and directionality of tree construction. In Lin, V., Krause, C., Bruening, B. & Arregi, K. (eds.), *Papers on morphology and syntax, cycle two* (MIT Working Papers in Linguistics **34**). Cambridge, MA: MITWPL. 67–105.
Richards, N. (2002). Movement in a top-down derivation. In Epstein, S. & Seely, D. (eds.), *Derivation and explanation in the Minimalist program*. Malden, MA: Blackwell.
Riecke, F., Warland, D., deRuyter van Steveninck, R. & Bialek, W. (1996). *Spikes: exploring the neural code*. Cambridge, MA: MIT Press.
Saito, M. (1985). *Some asymmetries in Japanese and their theoretical implications*. Ph.D. dissertation, MIT.
Schneider, D. (1999). *Parsing and incrementality*. Ph.D. dissertation, University of Delaware.
Selkirk, E. (1984). *Phonology and syntax*. Cambridge, MA: MIT Press.
Shastri, L. & Ajjanagadde, V. (1993). From simple association to systematic reasoning. *Behavioral and Brain Sciences* **16**. 417–494.

Spelke, E. (2003). What makes us smart? Core knowledge and natural language. In Gentner, D. & Goldin-Meadow, S. (eds.), *Language in mind: advances in the study of language and thought*. Cambridge, MA: MIT Press. 277–311.

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

Steedman, M. (2000b). Information structure and the syntax-phonology interface. *Linguistic Inquiry* **31**. 649–689.

Stowe, L. A. (1986). Parsing *WH*-constructions: evidence for on-line gap location. *Language and Cognitive Processes* **1**. 227–245.

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

Townsend, D. & Bever, T. (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. (2004). Neural blackboard architectures of combinatorial structures in cognition Ms., University of Leiden & Technische Universität München.

Wasow, T. (2002). *Postverbal behavior*. Stanford, CA: CSLI Publications.

Whitney, C. & Weinberg, A. (2003). Representing a parse in the brain: the TPARRSE model. Ms., University of Maryland.

Yoshida, M., Aoshima, S. & Phillips, C. (2004). Relative clause prediction in Japanese. Ms., University of Maryland.

Zec, D. & Inkelas, S. (1990). Prosodically constrained syntax. In Inkelas, S. & Zec, D. (eds.), *The phonology–syntax connection*. Chicago: The University of Chicago Press. 365–378.

Authors' address: Department of Linguistics, University of Maryland,
College Park, MD 20782, U.S.A.
E-mail: colin@umd.edu, ellenlau@mail.umd.edu