Routledge
Taylor & Francis Group

Some arguments and nonarguments for reductionist accounts of syntactic phenomena

Colin Phillips

Department of Linguistics, Program in Neuroscience and Cognitive Science, University of Maryland, College Park, MD, USA

Many syntactic phenomena have received competing accounts, either in terms of formal grammatical mechanisms, or in terms of independently motivated properties of language processing mechanisms ("reductionist" accounts). A variety of different types of argument have been put forward in efforts to distinguish these competing accounts. This article critically examines a number of arguments that have been offered as evidence in favour of formal or reductionist analyses, and concludes that some types of argument are more decisive than others. It argues that evidence from graded acceptability effects and from isomorphism between acceptability judgements and on-line comprehension profiles are less decisive. In contrast, clearer conclusions can be drawn from cases of overgeneration, where there is a discrepancy between acceptability judgements and the representations that are briefly constructed on-line, and from tests involving individual differences in cognitive capacity. Based on these arguments, the article concludes that a formal grammatical account is better supported in some domains, and that a reductionist account fares better in other domains. Phenomena discussed include island constraints, agreement attraction, constraints on anaphora, and comparatives.

Keywords: Syntax; Parsing; Overgeneration.

WHAT IS AT STAKE?

Everybody should want to see syntactic theories that are as simple as possible, and yet the facts of natural language appear to be rather complex. There are a number of routes to achieving greater simplicity. One route is to identify rules and generalisations that cover an increasingly broad range of empirical phenomena. Much fine work in linguistics follows this route. A second route, and one that has been well-trodden by linguists and psychologists alike, is to simply ignore or deny many of the rich phenomena that motivate the creation of complex theories. In some instances this may be the correct route, if the phenomena have been incorrectly characterised. A third route, and the one that I am concerned with here, is to propose that some of the phenomena that have been used to motivate complex syntactic theories are in fact consequences of other cognitive mechanisms, and hence do not need to be explained by formal syntactic theories. Accounts of this third type often argue that a class of sentence structures that speakers judge to be unacceptable are not, in fact, grammatically ill-formed, but instead are perceived as bad because they overburden the language processing system in some fashion. I will refer to such accounts as *reductionist* accounts of syntactic phenomena.

Across the past 40 years one finds reductionist accounts of many different types of syntactic phenomena. These include accounts of scope phenomena that attribute marked or unavailable scope interpretations to the operations of the parsing mechanism (Anderson, 2004; Reinhart, 2006); accounts of constraints on anaphora such as Binding Principle C (Chomsky, 1981) that attribute the constraints to effects of linear ordering (Gordon & Hendrick, 1997) or the pragmatics of discourse interpretation (Harris & Bates, 2002). They include accounts of superiority effects that attribute the constraints to independently motivated factors that cause comprehension difficulty (Arnon, Snider, Hofmeister, Jaeger, & Sag, in press), and accounts of partial agreement effects (Steiner, 2009) and the Right Roof Constraint (Grosu, 1973) that appeal to properties of on-line interpretation. There are "agreement attraction" phenomena that have been analysed both as formal grammatical phenomena (e.g., den Dikken, 2001; Kayne, 1989; Kimball & Aissen, 1971) and as errors that arise during language processing (e.g., Bock & Miller, 1991; Clifton, Frazier, & Deevy, 1999). Related controversies can also be found in the literature on child language development, where one finds competing grammatical and reductionist accounts of such phenomena as subject omission (Bloom, 1990; Hyams & Wexler, 1993), root infinitives (e.g., Phillips 1995/2010; Poeppel & Wexler, 1993; Rizzi, 1994), and constraints on anaphora (e.g., Grimshaw & Rosen, 1990; Grodzinsky & Reinhart, 1993; Hendriks & Spenader, 2005/2006; Thornton & Wexler, 1999). The domain that has attracted perhaps the largest number of reductionist

analyses is constraints on unbounded dependencies, so-called "island constraints" (Goodluck & Rochemont, 1992; Ross, 1967). So for example it is sometimes suggested that the unacceptability of *wh*-dependencies that cross subject boundaries, such as *What did the attempt to learn about ultimately confuse the students?,* is a consequence of parsing difficulty rather than formal grammatical constraints (Kluender, 2005; Pritchett, 1991). Many of these phenomena are discussed in more detail below.

Discussions over formal vs. reductionist accounts of syntactic phenomena often lead to exasperation on both sides, with advocates of either position perceiving that their arguments are unappreciated or even ignored. My aim in this article is to survey the different types of arguments that appear in these debates, with the possibly futile goal of tempering some of the acrimony that one encounters in such arguments. I have participated in debates over formal vs. reductionist accounts in a number of different syntactic domains, and I have found myself on different sides of the issue in different domains. I think that the reductionist position is the correct account for some phenomena, and is not the correct account for some other phenomena. This article aims to set out the reasons for this. But a number of preliminaries must be addressed before we proceed.

First, I will take it for granted here that there are at least some formal grammatical constraints, for without this there would be little to discuss. Most of the discussion will not depend on the specifics of any particular syntactic model, and the approaches that I group together under the heading of "formal" theories is broad enough to encompass such uncommon bedfellows as minimalist grammars (Chomsky, 1995) and probabilistic grammars (Booth & Thompson, 1973; Manning & Schütze, 1999). These approaches share the assumption that syntactic constraints are explicitly mentally represented, and hence they count for current purposes as formal theories, to be contrasted with theories in which the same constraints are assumed to be epiphenomenal.

Second, it is important to bear in mind that putative formal grammatical constraints could exert their effects in different ways. On the one hand, a constraint could apply as a *constraint on structure generation*, meaning that the constraint makes it impossible for speakers to construct mental representations that violate the constraint. On the other hand, a constraint could apply as a *filter* on generated structures, meaning that speakers are able to mentally represent structures that violate the constraint, but that they somehow also encode the fact that those representations are grammatically illicit. So, for example, English speakers have little difficulty in understanding agreement violations such as **The dog are barking* or argument structure anomalies such as **The official presented the winner a medal*, but they also easily recognise that the sentences are ill-formed. For this reason, it is

important to remember that evidence that a speaker is able to represent a given structure does not entail that the speaker treats that structure as well-formed.

Third, it is important to distinguish *reductionist* and *grounded* accounts of syntactic phenomena, although they are often grouped together under the label of "processing accounts". A reductionist account maintains that the phenomenon in question can be explained without any recourse to formal syntactic constraints, and that the phenomenon can be derived entirely from independent mechanisms. In contrast, a grounded account shares with formal accounts the assumption that the phenomenon reflects a syntactic constraint that is explicitly mentally represented. However, grounded accounts maintain that the existence of the constraint is itself motivated (historically, evolutionarily, or developmentally) by the fact that the constraint reduces language processing difficulty. For example, in the domain of resource-based accounts of constraints on unbounded dependencies, one encounters genuine reductionist accounts (Hofmeister & Sag, 2010; Kluender, 1998; Kluender & Kutas, 1993; Maratsos & Kowalsky, 2005) as well as grounded accounts of islands and constraints on nested dependencies (Berwick & Weinberg, 1984; Fodor, 1978; Hawkins, 1999). This distinction is important, as evidence against reductionist accounts of a given phenomenon may still be compatible with grounded accounts. Grounded accounts of syntactic constraints are not direct claims about the form of the constraints in the minds of speakers, but rather are claims about how the constraints arose in human language. They argue that certain types of constraints are more natural or more adaptive, because of their benefits for language processing. But unlike reductionist accounts, grounded accounts continue to maintain that the constraints are genuine formal restrictions. As such, they are extremely difficult to test empirically. I will have nothing further to say about grounded accounts here.

Fourth, the understanding of real-time comprehension and production mechanisms has advanced to the point where reductionist accounts of linguistic phenomena need to do more than appeal to all-purpose notions such as "processing difficulty". There are many different possible bottlenecks in comprehension and production, and theories should be explicit about this. A sentence may be difficult to understand because it is prone to misanalysis due to temporary ambiguity ("garden path" effects), or because of possible errors in retrieving the elements that participate in a syntactic dependency ("interference" effects). Alternatively, the sentence may have an overall size or complexity that exceeds the comprehender's capacity ("overload" effects). Or the sentence may be difficult to interpret because it overburdens the comprehender's ability to relate the linguistic form to a suitable mental model. Or a specific component of the sentence's representation may be difficult to construct incrementally, due to constraints on the left-to-right presentation of the sentence. Or the sentence may pose difficulties for

syntactic encoding in language production rather than for language comprehension. All of these and other mechanisms could potentially interfere with judgements of acceptability. Reductionist accounts that tie particular phenomena to specific real-time mechanisms are as explicit as their formal counterparts and they are eminently testable, whereas accounts that appeal to all-purpose notions of "processing difficulty" are as hard to test as they are vague.

Finally, it is important to distinguish genuine simplification of linguistic theories from spurious simplification that merely passes to others the burden of solving problems. If a syntactic constraint can be shown to be truly epiphenomenal, then this is potentially a good thing: it leads to simpler formal models and also reduces the burden on theories of language evolution. However, from the perspective of the language learner this simplification is of limited benefit. The child's task is to learn the properties of the ambient language, using input that is generated by language users and that is filtered through the child's own limited-capacity comprehension system. The child must somehow distinguish those patterns in the input that reflect properties of the grammar from those patterns that reflect constraints on language use, and the child's own immature comprehension system may lead to misanalyses that cloud the input data. The greater the range of factors that shape the sentences that the child encounters, the more acute is the child's problem in assigning blame for patterns in the input. I know of no existing account of how a learner could reliably tell the difference between those properties of the input that reflect the grammar of the target language and those that reflect constraints on language use. In this regard, reductionist accounts that reduce the burden on the syntactician do not automatically reduce the burden on the child.

With these preliminaries in mind we now turn to a series of different arguments that arise in discussions of formal vs. reductionist accounts of syntactic phenomena. I first discuss some arguments that I consider to be less effective in deciding between formal and reductionist accounts, and then move to arguments that are potentially more informative. All of the arguments that I discuss here have been used in past work as evidence for or against reductionist accounts, or have been construed as arguments for reductionist accounts even when they were not intended as such.

NONARGUMENT 1: USER-FRIENDLY GRAMMARS

A number of recent formal syntactic models have been proposed in which syntactic structures are assembled in a (roughly) left-to-right order, matching the order in which sentences are spoken and understood (e.g., Cann, Kempson, & Marten, 2005; Kempson, Meyer-Viol, & Gabbay, 2001; Phillips,

1996, 2003). These approaches take notions of derivational order more literally than is the norm in standard linguistic theories, which tend to explicitly distance themselves from claims about real-time operations. As such, models with left-to-right derivations aim to also describe important aspects of language processing. Yet it would be misleading to view these as reductionist accounts of syntactic phenomena, as they employ familiar formal devices and incorporate close analogues of the formal constraints found in other syntactic models. It is unfortunate, therefore, that linguists sometimes characterise these proposals as "processing accounts", grouping them together with genuine reductionist accounts of syntactic phenomena. This is an illustration of how unhelpful the cover-term "processing accounts" is, as it encourages conflation of sharply divergent claims about human language.

There is one narrow respect in which left-to-right syntactic models count as reductionist. These approaches typically argue that left-to-right derivations allow for simpler formal accounts of specific syntactic phenomena. For example, Phillips (1996, 2003) argues that left-to-right derivations can explain the conflicting outcomes of different constituency diagnostics, reducing the need for diagnostic-specific stipulations. Similarly, Shan and Barker (2006) argue that left-to-right interpretive mechanisms lead to a superior account of crossover phenomena, and Kempson and Kiaer (2010) propose that their *Dynamic Syntax* approach provides a simpler account of long-distance scrambling effects in Korean and Japanese. The same is true of many other proposals based on left-to-right structure building (Chung, Ladusaw, & McCloskey, in press; Fortmann, 2005; Guilliot, 2006; Richards, 1999; Schlenker, 1999). But all of these are reductionist proposals only in the narrow sense that they seek to simplify formal accounts of particular syntactic phenomena. In all other respects, these are standard formal syntactic models.

NONARGUMENT 2: ONE-DIMENSIONAL RATINGS

Acceptability ratings for sentences are graded, and this is particularly clear when speakers are asked to provide ratings using a continuous scale. Ratings of perceived difficulty are also graded. This means that when speakers are asked to judge sentences that vary in terms of both grammaticality and difficulty, their judgements tend to vary along the same scale. In on-line studies, reaction time is a similarly 1-dimensional measure that is impacted by difficulty and ill-formedness alike. The existence of gradedness is sometimes seen as evidence for a reductionist account of a syntactic phenomenon. The argument runs as follows: if the putative grammatical constraint affects ratings in the same way that manipulations of processing

difficulty do, then surely it is more parsimonious to assume that the phenomena have the same origin (Hofmeister & Sag, 2010).

This argument is straightforwardly dispatched. If speakers are asked to judge a sentence using a 1-dimensional scale, then it is unsurprising that they give judgements that make use of that scale. If the question under investigation is whether variation in acceptability judgements reflects one, two, or more orthogonal cognitive dimensions (e.g., difficulty vs. grammaticality), then we should not expect to resolve that question by asking speakers to give 1-dimensional ratings, and checking whether manipulation of an extra-grammatical factor modulates ratings in the same manner as the putative grammatical constraint. The same is true for 1-dimensional reading time variations. This is no different than the use of IQ tests as measures of cognitive abilities. The 1-dimensional nature of IQ scores could never confirm or refute the notion that intelligence reflects multiple independent cognitive dimensions. The widespread use of 1-dimensional scales to measure cognitive abilities makes it tempting to think of intelligence as a 1-dimensional property in humans. We now know that this is almost certainly incorrect, but it remains convenient to use the 1-dimensional scale for many practical purposes. By parity of reasoning, the gradedness of 1-dimensional linguistic acceptability judgements tells us as much about the contributors to acceptability as does the gradedness of 1-dimensional IQ scores.

NONARGUMENT 3: ISOMORPHISM

A more interesting case is the argument from *isomorphism*. Many experimental studies have asked whether real-time language processing is subject to the same syntactic constraints that are revealed through off-line acceptability judgements. Many such studies have revealed close alignment between on-line processes and off-line judgements. Many behavioural studies show that the same representations that are judged to be unacceptable in rating studies are avoided in on-line comprehension (e.g., Kazanina, Lau, Lieberman, Yoshida, & Phillips, 2007; Nicol & Swinney, 1989; Stowe, 1986; Sturt, 2003; Traxler & Pickering, 1996; Wagers & Phillips, 2009). Electrophysiological studies of syntactic anomaly typically show that unacceptable sentences elicit Event Related Potentials (ERP) responses such as the left-anterior negativity (LAN) or P600 components that begin within 300–600 ms of the anomalous word (Friederici, Pfeifer, & Hahne, 1993; Hagoort, Wassenaar, & Brown, 2003; Neville, Nicol, Barss, Forster, & Garrett, 1991; Osterhout & Mobley, 1995; for reviews see Friederici & Weissenborn, 2007; Kaan, 2007).

For example, in a series of studies led by Nina Kazanina and Ellen Lau (Kazanina et al., 2007) we tested whether on-line resolution of backwards anaphora is immediately sensitive to binding Principle C (Chomsky, 1981).

Pronouns typically follow their antecedents, but there are also cases where a pronoun may precede its antecedent (1a). However, there are configurations where this is not possible (1b), and this is commonly captured by Principle C, which requires that a pronoun not c-command its antecedent.[1]

(1) a.   While she$_i$ was reading the book, Mary$_i$ ate an apple.
    b.   *She$_i$ ate an apple while Mary$_i$ was reading the book.

Previous studies on the processing of backwards anaphora have shown that when comprehenders encounter a pronoun early in a sentence they actively search for an antecedent for the pronoun, anticipating positions for likely antecedents in advance of seeing them in the input. For example, van Gompel and Liversedge showed in an eye-tracking study that a *gender mismatch effect* is observed when comprehenders read a noun that mismatches in gender with a preceding pronoun (van Gompel & Liversedge, 2003). The gender mismatch effect at the underlined noun in sentences like (2) suggests that comprehenders expect that the main clause subject will be an antecedent for the pronoun, and are therefore surprised when it turns out to be incompatible with the pronoun.

(2) a   When he was fed up, the boy visited the girl very often.
    b.   When he was fed up, the girl visited the boy very often.

Kazanina and colleagues replicated the gender mismatch effect in a self-paced reading paradigm, but they also showed that no gender mismatch effect is observed on nouns that are excluded as potential antecedents for the pronoun, due to Principle C. This effect was consistent across three different structural environments that are all subject to Principle C (3a–c), and similar findings have been obtained in Japanese (Aoshima, Yoshida, & Phillips, 2009) and Russian (Kazanina & Phillips, 2010). The lack of gender mismatch effect suggests that comprehenders do not consider such positions to be candidate antecedent positions.

---

[1] There are at least two well-known classes of exception to Principle C in English, both of which arise in semantically well-defined contexts. The first class involves comparisons of multiple "guises" of the same individual (Reinhart, 1983; Heim, 1998), as in *He$_i$ then did what John$_i$ always did in such situations.* The second class of exceptions arises in cases where the embedded clause describes an event that interrupts the main clause event, as in *He$_i$ was threatening to leave when Billy$_i$ noticed that the computer had died* (Harris & Bates, 2002). Minor modifications to these examples re-introduce the ill-formedness of Principle C violations. Rather than undermining the validity of the Principle C constraint, such cases help to sharpen its formulation. For further discussion, see Kazanina (2005).

(3) a.   Because last semester she was taking classes full-time while
      <u>Kathryn</u>/<u>Russell</u> was working two jobs to pay the bills, Erica felt guilty.
      b.   It seemed worrisome to him that <u>John</u>/<u>Ruth</u> was gaining so much
      weight, but Matt didn't have the nerve to comment on it.
      c.   He/she chatted amiably with some fans while the talented young
      <u>quarterback</u> signed autographs for the kids, but Steve/Carol wished
      the children's charity event would end soon so he/she could go home.

Thus, in the case of Principle C, on-line interpretive processes exhibit
isomorphism with off-line judgements. Similar findings have been obtained
for other grammatical constraints, such as island constraints (McElree &
Griffith, 1998; Stowe, 1986; Traxler & Pickering, 1996; Wagers & Phillips,
2009; for review see Phillips & Wagers, 2007).

   Findings of isomorphism between on-line interpretation and off-line
judgements certainly invite the conclusion that the same mechanisms
are involved in the two domains. However, the conclusion that the same
mechanisms are involved tells us little about the specific nature of those
mechanisms. Isomorphism could indicate that formal constraints are
deployed in real-time processes, and is thus compatible with formal
grammatical accounts. Alternatively, isomorphism could indicate that limita-
tions on on-line sentence processing also impact off-line judgements, as
claimed by reductionist accounts. Both types of account capture such findings
equally well. For this reason, isomorphism effects cannot, in most cases, decide
among formal and reductionist accounts of grammatical phenomena.

   For purposes of distinguishing formal and reductionist accounts of
syntactic phenomena, more can be learned from cases where on-line processes
and off-line judgements diverge. We turn to these cases in Section 6.

NONARGUMENT 4: AMELIORATION

Another type of reductionist argument that one often encounters is the
argument from *amelioration*. Amelioration arguments typically start by
showing a sentence that is agreed to be unacceptable. Then the sentence is
changed by replacing some words, modifying the discourse context, or even
by slightly modifying the structure, in such a way that the new sentence is
clearly more acceptable. The argument then points to the acceptable
sentence, and takes this as evidence that the structure in question must be
grammatically well-formed after all, implying that even the initial unaccept-
able sentence is grammatically licit, and that its unacceptability must
therefore be due to extra-grammatical limitations.

   For example, extraction of *wh*-phrases from the complements of nouns is
typically degraded (4a), and this is often attributed to the *Complex NP*

Constraint (Ross, 1967), but it is also possible to create examples of extraction from NP complements that speakers find to be somewhat more acceptable (4b–c). Similarly, extraction from relative clauses is generally strongly unacceptable (5a), but there are also cases that sound far better (5b–c). It has been standard in syntactic theories of the past 40 years to assume that the (a) examples are genuinely ungrammatical and that the (b–c) examples are exceptional, due to additional syntactic and nonsyntactic factors. But an alternative approach is to take the (b–c) examples as evidence that the extractions are all really well-formed, and that the unacceptability of the (a) examples must be due to additional nonsyntactic factors (Deane, 1991; Hofmeister & Sag, 2010). Arguments from amelioration commonly feature in reductionist accounts of island phenomena.

(4) a.   ?*It was a new company that Simon spread the rumor that they started __.
b.   Which reports does the government prescribe the height of the lettering on __? (Ross, 1967)
c.   Nixon was one president that they had no trouble finding votes for the impeachment of __? (Deane, 1991)

(5) a.   *This is the book that I met a man who understands __?
b.   This is a paper that we need to find someone who understands __. (Chung & McCloskey, 1983)
c.   This is the child who there is nobody who is willing to accept __. (Kuno, 1976)

An argument from amelioration involves two assumptions. First, the argument assumes that acceptability entails grammatical well-formedness: this is the basis for the conclusion that the (b–c) examples above are grammatical. Second, the argument relies on the assumption that superficially similar sentences have the same syntactic structure: this is the basis of the inference that the grammaticality of the (b–c) examples above entails that the (a) examples are also grammatically well-formed. These are reasonable initial assumptions, but unfortunately there are a number of reasons to doubt both assumptions.

It is not difficult to find counterexamples to the assumption that acceptability entails grammaticality.

In naturally occurring speech we often encounter speech errors that pass relatively unnoticed, particularly in situations where the sentence is moderately complex and where the listener has a clear idea of what the speaker intends to say. A personal favourite is a couple of statements on climate change made by Republican governor Sarah Palin during her US vice presidential campaign in Fall 2008. In a CBS News interview, Palin told

reporter Katie Couric that "I'm not one to solely blame all of man's activities on changes in climate" (30 September 2008), and a few days later in a TV debate she said "I'm not one to attribute every activity of man to climate change" (2 October 2008). There are good scientific reasons to doubt the claim that Palin intended to convey, but the literal meaning of her statements is something that is almost trivially true. Both of these remarks were seen by huge TV audiences, and yet the errors were probably not noticed by the vast majority of viewers. Examples such as this cast doubt upon speakers' ability to consistently notice errors.

There are also cases of more systematically studied "grammatical illusions", sentences that are judged to be acceptable despite being clearly ill-formed. For example, speakers typically judge the clausal comparative in (6) to be acceptable upon first encounter, despite the fact that it is gibberish.

(6) *More people have been to Russia than I have.

The problem with examples like (6) is that the main clause subject *more people* implies a comparison of quantities of individuals, yet the comparative clause provides no quantity of individuals that can serve as the basis for comparison. It is therefore surprising that most speakers do not notice this problem on first encounter with such sentences. This illusion was first noted anecdotally by Montalbetti (1984), and is presented by Townsend and Bever (2001) as striking evidence for the first-pass heuristic analyser component in their architecture. In a series of acceptability, rating studies led by Alexis Wellwood we have recently demonstrated that the illusion is robust and can be manipulated (Wellwood, Pancheva, Hacquard, Fults, & Phillips, 2009). We have argued that this illusion involves a misinterpretation of the sentence as involving a comparison of quantities of events, and in support of this we have shown that the strength of the illusion varies according to whether the predicate is "repeatable". Ungrammatical comparatives with predicates that an individual can carry out repeatedly (7a) receive higher ratings than variants with predicates that an individual can carry out only once (7b).

(7) a.   *More undergrads call their families during the week than I do.
    b.   *More New Yorkers began law school this semester than I did.

Speakers do notice the illusion in (6) when it is pointed out explicitly, or after a few trials of judging similar sentences, but the typical initial reaction is that the sentences sound just fine. The construal of an individual comparison as an event comparison may be related to other cases where assertions about quantities of events are formulated as assertions about quantities of individuals (Barker, 1999; Krifka, 1990). For example, in (8) the intended claim involves a quantity of ticket purchases, and it would not be considered

to be falsified if it turned out that, in fact, the increased sales were due to a smaller number of purchasers. It remains unclear why speakers should so willingly entertain an event comparison interpretation in sentences like (6) that clearly do not permit such an interpretation, but illusions like (6) clearly indicate that judgements of acceptability do not entail well-formedness.

(8) More people bought lottery tickets last year than ever before.

There are cases where an even more transparently ungrammatical sentence receives surprisingly high ratings in controlled judgement studies. (9a) is a standard example of a hard-to-process double centre embedding. In (9b) one of the verbs has been removed. This ought to lead to a perception of severe ill-formedness, yet it appears that this is not the case (Frazier, 1985). Experimental tests have found that such sentences receive ratings that are equivalent to (Gibson & Thomas, 1999) or even higher than the well-formed counterpart in (9a) (Christiansen & MacDonald, 2009).

(9) a.   The patient the nurse the clinic had hired was treating met Jack.
    b.   The patient the nurse the clinic had hired met Jack.

So while it is almost certainly the case that ill-formed sentences are normally perceived as such by speakers, there is good evidence of uncontroversially ill-formed sentences that speakers judge as acceptable.

The second assumption of arguments from amelioration is that superficially similar sentences share the same syntactic structure. This is, of course, a reasonable starting hypothesis, but it must be treated with caution. A series of recent studies on cross-language variation in island constraints are particularly instructive in this regard, as they offer compelling evidence that what appear to be "acceptable island violations" are in fact instances of alternative structures that happen to resemble islands. Studies of apparent island-insensitivity in East Asian and Scandinavian languages paint a similar picture.

Various authors have pointed out that East Asian languages like Chinese, Japanese, and Korean allow extraction from relative clauses under certain circumstances, as shown in (10) (Huang, 1982; Kuno, 1973; Sohn, 1980). All of these appear to challenge the generality of the ban on extraction from relative clauses.

(10) a.   [$_{RC1}$ [[$_{RC2}$ __$_i$ __$_j$ ki-tei-ru] huku]-ga$_j$ yog-ore-tei-ru] sinsi$_i$
      (Japanese)
        wear-asp-pres clothes-nom dirty-mid-asp-pres gentleman
        "the gentleman that the clothes that (he) is wearing are dirty".
    b.   [$_{RC1}$ [[$_{RC2}$ __$_i$ __$_j$ kanjian de] xuesheng$_j$] lai-le de] na-ge ren$_i$
      (Chinese)

see de student come-asp de that-cl person
"the man that the student who (he) saw came"

c.   [$_{RC1}$ [$_{RC2}$ __$_i$ __$_j$ cohaha-nun] kangaci-ka$_j$ cwuk-un] ai$_i$ (Korean)
like-adn dog-nom die-adn kid
"the kid who the dog which [he] liked died."

However, recent studies in each of Korean (Han & Kim, 2004), Chinese (Hsu, 2006), and Japanese (Sakai, 1994; Ishizuka, 2009) have argued that examples like (10) are only apparent violations. In each language the authors argue that examples like (10) do not involve extraction out of a relative clause, but rather are instances of merely local relativisation of a higher NP that is known in the Japanese–Korean literature as the *major subject* in a double nominative construction. They argue that the major subject in turn controls a null subject in the most deeply embedded relative clause. In this way, independently motivated properties of the languages in question yield the surface appearance of extraction from a relative clause, without such extractions actually occurring. In support of this, the authors demonstrate a series of parallels between constraints on the distribution of major subject constructions and constraints on "double relative clauses" like (10). To take just one example, double nominative constructions in Korean cannot be formed with transitive predicates. Accordingly, apparent extraction from a relative clause is also impossible when the lower relative clause occupies the object position of a transitive predicate (11) (Han & Kim, 2004). Sentences like (11) cannot be parsed as major subject constructions, and hence they require a standard relative clause extraction analysis, leading to the same unacceptability seen in their English counterparts.

(11) *[$_{RC1}$ wuli pan haksayng-i [$_{RC2}$ ip-un] yangpok-ul po-n] sinsa
our class student-nom wear-adn suit-acc see-adn gentleman
"the gentleman who a student from our class saw the suit which [he] wore"

These arguments suggest that relative clauses block extraction in Korean, Chinese, and Japanese, just as in English, and hence it would be misleading to infer from examples like (10) that extractions from relative clauses are well-formed in general in these languages.

Similar arguments have been presented for well-known cases of acceptable extraction from relative clauses in various Scandinavian languages (Allwood, 1982; Engdahl, 1982; Erteschik-Shir, 1973; Nordgaard, 1985; Taraldsen, 1982). Examples like (12a) are judged as perfectly acceptable by speakers of these languages, and they are reported to be relatively easy to find in corpora of naturally occurring speech (Engdahl, 1997). However, their distribution is constrained, and many other cases of extraction from relative clauses in these

languages are just as unacceptable as their English counterparts (Allwood, 1982; Engdahl, 1997). In a recent study, Kush argues that the cases of acceptable extractions consistently show properties of small clause constructions in these languages, and hence that examples like (12) involve extraction from a small clause rather than a relative clause (Kush, 2010). Kush shows that when such examples are modified in a variety of ways to block construal as a small clause, such as by using a different main clause verb (12b), then standard constraints on extraction from relative clauses re-emerge.

(12) a.   Den teorin känner jag ingen som tror på.(Swedish)
           That theory know I nobody that believes in
           "That theory, I know nobody who believes in".
      b.   *Den här teori, finns det ingen som tror på.
           "*That theory, there is nobody who believes in*".

These arguments about apparent cross-language variation in island constraints serve as a useful reminder of the problem of structural ambiguity in language, and they further undermine the force of arguments from amelioration. Returning to the examples of amelioration in English in (4–5), the more acceptable examples (4b–c, 5b–c) may indicate that the less-acceptable-but-superficially-similar examples (4a, 5a) are also well-formed. But they may instead reflect the fact that the examples have different structures.

I should emphasise, though, that this does not mean that amelioration effects are uninformative or should be discounted. Nor do these examples show that amelioration effects always reflect structural ambiguity. On the contrary, efforts to understand such effects consistently yield valuable insights, and the correct account of each case must be assessed on a case-by-case basis. Most likely, some of the cases where unacceptable sentences can be ameliorated reflect structural ambiguity, some cases reflect grammatical illusions, and some genuinely show that the unacceptable sentences are structurally well-formed. What the examples in this section show is that amelioration effects do not immediately reveal the underlying nature of unacceptability effects. Rather, they serve as useful starting points for further investigation.

ARGUMENT 1: OVERGENERATION

I have argued that we learn little about the nature of constraints from instances of isomorphism, that is, situations where on-line language processing mechanisms construct only representations that are judged as acceptable in off-line tasks. Such situations are less decisive because isomorphism could reflect the on-line and off-line impact of grammatical constraints, or it could reflect parallel on-line and off-line effects of resource

limitations. In contrast, it is particularly informative when we encounter situations where on-line processes construct representations that are judged as unacceptable in off-line tasks.

A priori we might expect isomorphism to be pervasive: human language would make available a representation-building system, and that would define all and only the representations that are constructed on-line, which in turn would correspond perfectly with the representations that speakers judge to be acceptable. But isomorphism is not so pervasive. Speakers are able to represent many types of sentences that are uncontroversially ill-formed, such as subject–verb agreement violations ("They is here") and argument structure violations (e.g., *"She mentioned him her concerns"). For practical purposes, this ability is useful, as it allows comprehension of ungrammatical input, as occurs in speech errors or in conversation with non-native speakers. Meanwhile, for the current concerns the ability to represent ungrammatical input is both a help and a hindrance. On the one hand, it implies that a speaker's ability to understand or represent a given sentence cannot be used as evidence that the sentence is grammatical for the speaker. On the other hand, discrepancies between what speakers are able to represent and what they judge to be acceptable can provide a powerful tool for understanding the nature of linguistic constraints.

Situations where there is a discrepancy between broader and narrower sets of representations that speakers generate under different circumstances can be characterised as instances of *overgeneration*. The narrower set consists of those representations that speakers allow under any circumstances—I refer to these as the *stable* forms. The broader set also includes those representations that speakers allow only under specific circumstances—I refer to these additional representations as the *overgenerated* forms. The difference between the stable forms and the overgenerated forms might be seen, for example, by comparing what speakers do in on-line comprehension with the results of untimed judgement studies with the same speakers. Alternatively, the difference may be seen in comparing the representations that individual speakers entertain in different pragmatic contexts. But whatever the cause of the overgeneration, the discrepancy between the stable and the overgenerated forms can be extremely informative in attempts to understand the nature of linguistic constraints.

The critical point is this: if the overgenerated forms are judged as unacceptable or are unavailable under some circumstances, despite the fact that speakers are clearly able to construct those representations in some situations, then we can draw clearer conclusions about the nature of the constraint that blocks the overgenerated forms under normal circumstances. In particular, it is hard in such cases to argue that the overgenerated forms are ruled out by limitations of the human representation-building capacity, since we have evidence that speakers are able to construct exactly those representations. It therefore becomes more likely that the constraint that

normally blocks the overgenerated forms is a formal grammatical constraint. As such, I think that this *argument from overgeneration* offers one of the better forms of evidence in favour of formal grammatical analyses of linguistic phenomena.

In some cases we are primarily interested in the nature of the constraint that typically blocks the overgenerated forms, and in such cases the argument just outlined can provide evidence for a formal grammatical account of the constraint. But in other cases we are primarily interested in the status of the overgenerated forms themselves. If the overgenerated forms are normally ruled out by a grammatical constraint, then it is plausible that they are grammatically ill-formed, and that they either reflect the temporary failure of grammatical structure building mechanisms, or they reflect intermediate steps in the construction of grammatically well-formed representations.

With these preliminaries in mind, we now turn to some examples of overgeneration and how it helps to diagnose the nature of constraints.

Island constraints

One case where overgeneration may provide useful evidence on the nature of a constraint comes from work on subject island effects in language processing. It has long been known that sentences with extraction from complex subjects are judged unacceptable (13), and this can be captured by a formal constraint that bans such dependencies (Cattell, 1976; Chomsky, 1973; Huang, 1982; Kayne, 1983). Alternatively, the unacceptability could reflect limitations that apply to on-line structure building (Kluender, 2005; Pritchett, 1991).

(13) *What did the attempt to repair ___ ultimately damage the car?

Reductionist accounts of island constraints make a straightforward prediction: if an island constraint is ultimately a consequence of limits on on-line structure building, then we should not find that the constraint is spontaneously violated during on-line structure building. Many of the previous studies of the on-line effects of island constraints have tested whether speakers construct illicit filler–gap dependencies into complex subjects during real-time comprehension, and a consistent finding has been that they do not (Omaki & Schulz, in press; Stowe, 1986; Traxler & Pickering, 1996). Such findings are compatible with reductionist accounts, but since they are findings of isomorphism they are, of course, equally compatible with formal grammatical accounts of the constraint. A more informative case, for current purposes, is the finding that comprehenders do appear to construct filler–gap dependencies into complex subjects under certain restricted circumstances (Phillips, 2006).

Although filler–gap dependencies that span a subject NP boundary are typically judged to be strongly unacceptable (14a), they are possible in so-called parasitic gap constructions, as illustrated in (14c), where the filler is associated with two gaps, one inside a subject and another in object position. The gap inside the subject is said to be "parasitic" upon the object gap because its well-formedness depends on the presence of the object gap. In contrast, the object gap is fully acceptable on its own (14b). Parasitic gap phenomena are very interesting for formal linguistic models, since they raise the challenge of explaining how the combination of a licit gap with an otherwise illicit gap yields an acceptable result (Engdahl, 1983; Kayne, 1983; for a useful anthology see Culicover & Postal, 2001). But they are potentially even more informative for purposes of adjudicating between grammatical and reductionist accounts of island phenomena.

(14) a.   *Which parts did [the attempt to repair ___] ultimately damage the car?
     b.   Which parts did [the attempt to repair the car] ultimately damage ___?
     c.   Which parts did [the attempt to repair ___] ultimately damage ___?

Phillips showed in a self-paced reading study that comprehenders attempt to construct a filler–gap dependency at the point of the verb inside the complex subject NP. Evidence for this came from manipulation of the plausibility of the semantic relation between the wh-phrase and the embedded verb, which led to reading time slowdowns at the embedded verb. None of the experimental materials actually contained gaps inside the embedded subject, so the reading time slowdowns reflected actions taken by the parser in advance of bottom-up evidence about the position of the gap. A separate acceptability judgement task confirmed the judgements shown in (14). These findings together constitute a case of temporary overgeneration. The gap inside a subject in (14a) is reliably judged to be unacceptable, yet during on-line sentence comprehension speakers appear to actively posit a gap in exactly that position. This action is likely related to the possibility that the sentence may turn out to be a parasitic gap construction, and in support of this Phillips shows that other complex subjects that do not allow parasitic gaps do not induce active gap creation, but this point is not crucial to the current argument. Since speakers are clearly able to construct filler–gap dependencies that span a subject NP boundary, it should be difficult to maintain a reductionist account of subject islands that attributes the unacceptability of (13) to limitations of on-line sentence comprehension mechanisms.

It should be noted, however, that this argument from overgeneration only indirectly establishes a link between the unacceptability of subject island violations (13) and formal grammatical constraints. What the argument does is challenge any account of subject island effects that appeals to limitations on on-line mechanisms. If the only alternative (or the most plausible alternative) to such accounts involves a formal grammatical constraint, then the formal grammatical account is supported. But if there are other feasible nonformal accounts of the unacceptability of subject islands then they could potentially be reconciled with the on-line findings.

Principle B in children

The argument from overgeneration may also be useful in understanding the nature of children's grammatical abilities, as can be shown using an example from the study of anaphora.

According to well-known formulations of grammatical constraints on anaphora, a pronoun may not have a clausemate antecedent. This means that a sentence like *John washed him* should not normally allow an interpretation where it means that John washed himself. There are different formal implementations of this restriction, which typically falls under the rubric of binding Principle B (e.g., Chomsky, 1981; Reinhart & Reuland, 1993). There is much more to be said about the syntactic, semantic, and pragmatic details of Principle B effects (e.g., Elbourne, 2008; Evans, 1980; Grodzinsky & Reinhart, 1993; Heim, 1998), but the simple formulation suffices for current purposes.

In the language acquisition literature there are many reports of children allowing interpretations of pronouns that violate Principle B. For example, Chien and Wexler (1990) showed in a large-scale series of picture verification studies with over 500 children aged 3–6 years that children frequently assented to sentences like *Is Mama Bear touching her?* when viewing a picture in which Mama Bear is touching herself. Similar findings have been obtained using diverse experimental measures and different languages, with over 30 published studies on the topic (e.g., Jakubowicz, 1984; McDaniel, Cairns, & Hsu, 1990; Spenader, Smits, & Hendriks, 2009; Thornton & Wexler, 1999; for review see Conroy, Takahashi, Lidz, & Phillips, 2009). Discussion of the phenomenon has focused on the question of whether children's errors reflect deficient grammatical knowledge (e.g., Hendriks & Spenader, 2005/2006; Thornton & Wexler, 1999) or problems in the implementation of an otherwise adult-like system of grammatical constraints (e.g., Conroy et al., 2009; Grimshaw & Rosen, 1990; Grodzinsky & Reinhart, 1993).

One aspect of children's behaviour that may provide important clues to the source of their errors is the substantial variability that is observed across studies. In studies with English-speaking children one finds rates of

nonadult-like pronoun interpretations that are as low as 10–15% in some studies (Conroy et al., 2009; Kaufman, 1988), and as high as 70% or greater in others (Matsuoka, 1997; McKee, 1992). In addition, whereas most comprehension studies elicit substantial numbers of nonadult-like interpretations from children, corresponding errors are not observed in children's spontaneous speech (Bloom, Barss, Nicol, & Conway, 1994; De Villiers, Cahillane, & Altreuter, 2006; Spenader et al., 2009). Conroy and colleagues argue that this variability is systematic, and that rates of nonadult-like interpretations in Principle B studies are somewhat predictable, based on the narrative contexts in which children are asked to judge the target sentences (see also Elbourne, 2005). In addition, Conroy and colleagues show that by modifying specific features of the test stories in a truth value judgement task they were able to significantly increase or decrease children's rate of nonadult-like interpretations. In studies that provided strong discourse support for both deictic and co-referential interpretations of the pronoun, they found a sharp contrast in judgements between sentences that are subject to Principle B (15a) and sentences that are exempt from Principle B (15b), using identical test scenarios. These judgements closely match adult judgements for the same items, and the contrast between judgements of (15a) and (15b) shows that the low acceptance of coreference in (15a) is unlikely to be due to the inaccessibilty of the co-referential interpretation, since that interpretation was readily chosen in (15b). In a follow-up study that altered the accessibility of the adult-like deictic interpretation, children's rate of errors for sentences like (15a) rose to 56%.

(15) a.   Grumpy painted him. *accept coref.*: 11% (children), 5% (adults)
     b.   Grumpy painted his costume. *accept coref.*: 80% (children), 83% (adults)

This pattern of performance in children represents an instance of overgeneration. Children consistently allow the same interpretations of pronouns that adults allow, and under certain circumstances they also allow interpretations that go beyond what adults normally allow. The contrast in children's judgements in (15a–b) strongly suggests that their grammars include a Principle B constraint. (The findings cannot, of course, decide between competing syntactic, semantic, and pragmatic accounts of Principle B in the theoretical literature.) Meanwhile, we suggest that children's "overgenerated" interpretations reflect difficulties in implementing the same mechanisms that adults use to apply Principle B in real-time interpretation. We propose that Principle B is not the kind of constraint that blocks the generation of illicit interpretations, but rather acts as a filter on candidate interpretations once they have been generated. When comprehenders encounter a pronoun they may attempt to link the pronoun with any feature-compatible antecedent in

the current sentence or discourse model, and then Principle B serves to inhibit any resulting interpretations that involve local binding. Evidence from adult on-line language processing suggests that adults might use a mechanism of this kind to implement Principle B. Although adults reliably reject local co-reference interpretations of pronouns in off-line judgement tasks, there is variability across studies in whether adults fleetingly consider local antecedents for pronouns (Badecker & Straub, 2002; Clifton, Kennison, & Albrecht, 1997; Hendriks, Banga, van Rij, Cannizzaro, & Hoeks, 2010; Kennison, 2003; Lee & Williams, 2006; Nicol & Swinney, 1989; Runner, Sussman, & Tanenhaus, 2006), mirroring the variability that we find in off-line judgement tasks in children. It is independently known that children have difficulty in inhibiting interpretations once they have been constructed during on-line comprehension (Mazuka, Jincho, & Oishi, 2009; Trueswell, Sekerina, Hill, & Logrip, 1999). Therefore, if the variability observed across studies in children's nonadult-like interpretation of pronouns is correlated with the relative accessibility of the licit and illicit interpretation of the pronoun, as we have argued, then children's errors are plausibly due to their difficulty in inhibiting interpretations.

Thus, in the case of Principle B we propose that the facts support a reductionist account of children's nonadult-like pronoun interpretations. Variability in children's performance in different situations suggests that this is a case of overgeneration, and the variability may be explained by a combination of the same generate-and-filter mechanisms that are used to implement Principle B in adults, in conjunction with independently known limitations in children's executive control capacities.

Agreement attraction

Another case where we encounter controversy over the status of over-generated forms can be found in the domain of agreement. As noted by Kimball and Aissen (1971) speakers sometimes accept sentences in which a relative clause verb agrees with the plural head of the relative clause (16a) rather than with the subject of the relative clause, as it should (16b). The same speakers reject variants of these sentence in which there is no plural noun for the relative clause verb to agree with (16c), showing that the judgements do not simply reflect insensitivity to agreement.

(16) a. (*)The drivers who the runner wave to each morning honk back cheerfully.
      b. The drivers who the runner waves to each morning honk back cheerfully.
      c.* The driver who the runner wave to each morning honks back cheerfully.

In the theoretical linguistics literature a number of studies have analysed cases of nonstandard agreement like (16a) as grammatically well-formed options in a substandard dialect (den Dikken, 2001; Kayne, 1989; Kimball & Aissen, 1971). For example, Kimball and Aissen originally presented the forms as features of an Eastern New England dialect, and others have noted that some speakers accept the forms whereas most reject them, and have taken this as direct evidence for grammatical differences between individual speakers. In contrast, in the psycholinguistics literature a number of studies have analysed such forms as errors that arise in the production or comprehension of standard subject–verb agreement relations (Bock & Miller, 1991; Clifton et al., 1999; Franck, Lassi, Frauenfelder, & Rizzi, 2006; Staub, 2009; Wagers, Lau, & Phillips, 2009). In contrast to the literature on island constraints, where there has been a lively debate between advocates of formal and reductionist accounts, these contrasting approaches to nonstandard agreement have been pursued largely in parallel, with little debate between proponents of the two broad approaches. Nevertheless, this does not make the phenomenon any less relevant to the current discussion.

In recent work led by Matt Wagers and Ellen Lau we have argued that nonstandard agreement forms like (16a) are consequences of fallible retrieval mechanisms that all speakers share (Wagers et al., 2009; see also Badecker & Lewis, 2008). In effect, our claim is that these forms are instances of overgeneration, and that the overgenerated forms are best captured by a reductionist account rather than a formal syntactic account. This stands in contrast to our conclusion in the case of island constraints.

The first step in the argument involves evidence that agreement forms like (16a) arise in all speakers. This runs counter to the view that such forms are the product of a nonstandard grammatical dialect found in a subset of speakers. That nonstandard dialect view is consistent with the results of informal polling, which typically shows that most speakers reject such forms but a non-negligible minority considers them acceptable. We obtain similar results in untimed pencil-and-paper acceptability rating tasks. However, we suspect that the judgements are misleading in this case. In on-line studies using self-paced reading or speeded acceptability measures we consistently find that an overwhelming majority of participants treat forms like (16a) as acceptable, as reflected in faster reading times or increased acceptance rates. Moreover, the minority of participants who do not show sensitivity to nonstandard agreement may simply reflect the noisiness of individual data in on-line studies. Our on-line measures have found no evidence that forms like (16a) are confined to specific dialects of English. Furthermore, very similar effects have been found in a speeded acceptability task in Brazilian Portuguese (Alcocer, França, Maia, & Phillips, 2010) and in sentence production studies in French (Franck et al., 2006), suggesting that subject–verb agreement processing is affected in similar ways across languages. We therefore propose

that all speakers share on-line agreement processing mechanisms that are susceptible to errors like (16a), and that the divergent judgements found in untimed tasks reflect differences among speakers in the ability to consciously access the fleeting products of the agreement processing system.

If it is correct that all speakers show at least fleeting acceptance of forms like (16a), but that most speakers reject the same forms in untimed judgement tasks, then this is another case of overgeneration. The stable forms conform to the standard subject–verb agreement pattern, and there is little real controversy that this pattern is the result of a formal grammatical rule. In this case the controversy surrounds the status of overgenerated forms like (16a). We have identified a number of reasons to favour a reductionist account that treats such forms as the product of noisy retrieval mechanisms used in on-line agreement processing, i.e., mechanisms that sometimes retrieve featurally appropriate agreement controllers from structurally inappropriate positions, rather than as the product of nonstandard grammatical agreement rules.

First, speakers are susceptible to agreement errors like (16a) in a probabilistic fashion. Sentences like (16a) consistently yield acceptance rates and reading-time distributions that are intermediate between unambiguously grammatical (16b) and unambiguously ungrammatical (16c). This is consistent with the notion that (16a) is the product of a noisy on-line mechanism.

Second, we consistently find a *grammatical asymmetry* in sensitivity to agreement between a relative clause verb and the head of the relative clause. When the standard subject–verb agreement relation is illicit, the head of the relative clause, which in the agreement processing literature is termed the *attractor*, gives rise to a (temporary) perception of well-formedness. But the opposite does not occur, i.e., when the standard subject–verb agreement relation is well-formed, a mismatch between the verb and the relative clause head does not reduce acceptability: manipulation of the features of the attractor has no effect in comparisons of well-formed sentences like (17a–b). In other words, agreement attraction impacts the processing of ungrammatical sentences, but not the processing of grammatical sentences. This asymmetry is unexpected if the grammar simply makes available the possibility of agreement relations between a relative clause verb and the head of the relative clause, but it receives a natural explanation in an account where the parser attempts to retrieve plural nominal features from its memory encoding of the sentence only in situations where it fails to encounter the grammatically expected agreement features on the verb. The retrieval of plural features from structurally inappropriate positions is straightforwardly captured by a language processing architecture such as proposed by Lewis and Vasishth and their colleagues (Lewis & Vasishth,

2005; Lewis, Vasishth, & van Dyke, 2006), which employs parallel cue-based retrieval mechanisms in content addressable memory.

(17) a.   The driver who the runner waves to each morning honks back cheerfully.
b.   The drivers who the runner waves to each morning honks back cheerfully.

Third, agreement attraction effects in production and comprehension alike exhibit a well-documented markedness contrast. Plural nouns induce agreement attraction, but singular nouns do not (Eberhard, Cutting, & Bock, 2005; Hartsuiker, Schriefers, Bock, & Kikstra, 2003; Vigliocco, Butterworth, & Garrett, 1996; Wagers et al., 2009). This contrast is readily captured in a retrieval-based account, and related effects are found independent of language, in visual search tasks, where a stimulus that includes a unique feature "pops-out" from an array of stimuli that lack that feature, but a stimulus that lacks a feature present in all others does not show pop-out effects (Treisman & Gelade, 1980). The markedness contrast in agreement is less expected in a grammatical account of nonstandard agreement.

In support of a formal grammatical account of agreement attraction effects Kayne (1989) and den Dikken (2001) point out that agreement attraction effects are induced by full NPs but not by pronouns, as in the contrast in (18). The argument relies on the assumption that grammars easily (and frequently) draw distinctions between pronouns and full NPs, and that such contrasts are less expected under an account that relies on noisy language processing mechanisms. In fact, further experimental evidence suggests that the lack of attraction effects from the pronoun in (18b) may be a consequence of lexical ambiguity in retrieval. In studies on agreement attraction in Dutch Hartsuiker and colleagues have found that object pronouns in subject–object–verb sentences do induce agreement attraction errors, but only when the object pronoun is case ambiguous. Case unambiguous pronouns do not induce attraction effects (Hartsuiker, Antón-Méndez, & van Zee, 2001). This suggests that the contrast in (18a–b) might reflect the case marking on English pronouns, rather than a grammatical contrast between pronouns and full NPs.

(18) a.   The identity of the participants are to remain a secret.
b.   The identity of them are to remain a secret.

In sum, as the examples in this section hopefully illustrate, evidence from instances of overgeneration can be particularly informative in distinguishing between formal and reductionist accounts of linguistic phenomena. When we encounter situations where there is a discrepancy between a narrower set of

"stable" forms that speakers accept under any circumstances and a broader
set of "overgenerated" forms that the same speakers accept only under
limited circumstances a couple of conclusions are likely. First, the division
between the stable forms and the overgenerated forms probably reflects a
formal grammatical constraint. Second, the overgenerated forms can arise
for a number of different reasons. They may reflect the parser's recognition
of a grammatically licit continuation of the sentence (parasitic gaps), or the
incomplete implementation of a grammatical constraint that applies in a
generate-and-filter fashion (Principle B in children), or it may reflect the use
of a grammatically appropriate retrieval operation in a noisy memory
architecture (agreement attraction).

ARGUMENT 2: CAPACITY EFFECTS

Re-interpretations of formal grammatical constraints in terms of constraints
on real-time comprehension mechanisms make an interesting prediction
about individual differences among speakers. If a representation is unac-
ceptable simply because it exceeds the resources that are available for
sentence parsing, then that representation should become more acceptable if
more resources are available. It is difficult to reliably expand parsing
resources within individuals, and therefore the most straightforward test of
this prediction is to compare individuals with different amounts of
processing resources. For example, a theory that attributes the unaccept-
ability of extraction from islands to resource overloads should predict that
speakers who have a sufficiently large pool of sentence processing resources
should find island extractions acceptable.

Despite the long-history of resource-based accounts of island constraints,
this prediction has not, to my knowledge, been put to the test until quite
recently, and the findings so far are not encouraging for reductionist
theories. Studies led by Sprouse in English and by Tokimoto in Japanese
indicate that speaker groups that vary in memory span are equally sensitive
to island constraints (Sprouse, Wagers, & Phillips, 2010; Tokimoto, 2009).
Sprouse and colleagues tested large groups of speakers (over 300 in total) in
a pair of studies that examined relations between sentence acceptability
ratings and measures of verbal working memory capacity. The studies used a
verbal memory-span task that required participants to recall lists of words
in a specific order (Cowan, 2000) and an *n-back* memory task that required
participants to judge, for varying values of *n*, whether the current word also
appeared *n* positions earlier in the list (Cohen et al., 1997; Poulton, 1954).
Both studies used an acceptability rating task that factorially manipulated
the presence of wh-extraction and four types of island contexts (*wh*-islands,

complex NP constraint, subject islands, and adjunct islands). One study used a Likert scale measure, the other used magnitude estimation (Bard, Robertson, & Sorace, 1996; Stevens, 1957). In both studies the acceptability ratings showed standard effects of island constraints: extraction consistently lowered ratings, the presence of island contexts without extraction lowered ratings to a minimal extent, and the combined effect of extraction and an island context was superadditive, with substantially lower ratings than were predicted by the combination of the two individual factors. In formal accounts of islands the superadditivity is attributed to violation of a grammatical constraint. In resource-based accounts of island effects the superadditivity is attributed to overload of comprehension resources. Resource-based accounts would receive striking confirmation if the super-additivity was reduced in individuals with greater resources. However, we found that the effect of island violations on acceptability ratings was consistent across speakers: the superadditive effect of combining extraction and an island context showed no correlation with memory span, irrespective of rating measure, memory-span measure, and island type. The only observed effect of memory span was that high span participants were more likely, in some conditions, to give *lower* ratings for syntactically complex sentences without island violations. Tokimoto (2009) reports similar findings in Japanese, using a yes/no acceptability rating task: high-span participants were not more likely to accept island violations, although they were slightly more likely to correctly accept control sentences in the study.

These results lend no support to resource-based accounts of island phenomena. They do not directly implicate formal accounts of island constraints, although the results are exactly what such accounts would predict. The results are also consistent with alternative reductionist accounts of islands that attribute the constraints to the semantic or functional "unnaturalness" of island violations (Deane, 1991; Erteschik-Shir, 1973).

One might object that studies of memory span and islands that focus on university students are poor tests of the resource-based account of islands, since they focus on a population that is too narrow. However, this objection amounts to a claim that the resource-based account of islands is not testable. The resource-based account predicts that the superadditive effect of islands on acceptability should be attenuated or eliminated in individuals with sufficiently large memory capacities. This attenuation is not found in the higher-performing group of university students, and nor is there any evidence that island effects are milder in professional linguists, who presumably have higher-than-average verbal processing abilities. Thus one might argue that island constraint violations tax human sentence comprehension capacity

to such an extreme degree that the resource-based account could only be tested in individuals with super-human verbal working-memory capacity. But that would lead to a further uncomfortable problem, due to the existence of cross-language variation in island effects. If island violations are so difficult in some languages that they exceed the sentence processing resources of any normal human being, then it is surely surprising that rather similar sentences should be judged acceptable in other languages.

Whereas island effects do not correlate with variation in memory capacity, we predict that correlations should obtain in those domains where we have argued that a reductionist account is correct. For example, we predict that children's propensity to interpret sentences that violate Principle B should correlate with their ability to inhibit and reanalyse incorrect parses. Also, we predict that individual adults who show greater susceptibility to agreement interference might show greater susceptibility to interference in other recall or visual search tasks. Very interesting findings from Kaan and colleagues suggest that this may indeed be the case (Kaan, Choi, Rowlings, Ballantyne, & Wijnen, 2009).

CONCLUSIONS

There is no question that it is important to determine what linguistic phenomena reflect formal grammatical constraints and what phenomena are reducible to independently motivated properties of language processing mechanisms. But it is typically not easy to distinguish competing formal and reductionist accounts of the same phenomena. I have attempted here to survey some of the arguments that have been used to argue for formal or reductionist accounts of various linguistic phenomena, or that are often understood to argue for reductionist accounts, even if they are not intended thus. In each case I have examined the assumptions that the argument. Some arguments that might be used in support of reductionist accounts, such as arguments from graded acceptability and amelioration, and arguments from isomorphism, turn out to be inconclusive, as they do not distinguish the predictions of formal and reductionist accounts. In contrast, I argued that more decisive evidence can be found in arguments from overgeneration, and arguments from individual differences in cognitive resources. Based on these arguments, I concluded that for some phenomena the formal grammatical account is more appropriate, and for other phenomena a reductionist account is better supported. Each case must be assessed on its merits.

REFERENCES

Alcocer, P., França, A., Maia, M., & Phillips, C. (2010, March). *Structure sensitive and insensitive retrieval of subjects in Brazilian Portuguese*. Paper presected at the 23rd annual CUNY Sentence Processing Conference, New York, NY.

Allwood, J. (1982). The complex NP constraint in Swedish. In E. Engdahl & E. Ejerhed (Eds.), *Readings on unbounded dependencies in Scandinavian languages* (pp. 15–32). Stockholm, Sweden: Almqvist & Wiksell.

Anderson, C. (2004). *The structure and real-time comprehension of quantifier scope ambiguity* (PhD dissertation). Northwestern University, Evanston, IL.

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

Arnon, I., Snider, N., Hofmeister, P., Jaeger, T. F., & Sag, I. A. (in press). Cross-linguistic variation in a processing account: The case of multiple wh-questions. In *Proceedings of the 32nd annual meeting of the Berkeley Linguistics Society*. Berkeley, CA: Berkeley Linguistics Society.

Badecker, W., & Lewis, R. (2008, March). *Agreement and sentence formulation: The role of working memory (retrievals) in language production*. Paper presented at the 21st annual CUNY Sentence Processing Conference, Chapel Hill, NC.

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.

Bard, E., Robertson, D., & Sorace, A. (1996). Magnitude estimation of linguistic acceptability. *Language*, *72*, 32–68.

Barker, C. (1999). Individuation and quantification. *Linguistic Inquiry*, *30*, 683–691.

Berwick, R., & Weinberg, A. S. (1984). *The grammatical basis of linguistic performance*. Cambridge, MA: MIT Press.

Bloom, P. (1990). Subjectless sentences in child language. *Linguistic Inquiry*, *21*, 491–504.

Bloom, P., Barss, A., Nicol, J., & Conway, L. (1994). Children's knowledge of binding and coreference: Evidence from spontaneous speech. *Language*, *70*, 53–71.

Bock, J. K., & Miller, C. A. (1991). Broken agreement. *Cognitive Psychology*, *23*, 45–93.

Booth, T. L., & Thompson, R. A. (1973). Applying probability measures to abstract languages. *IEEE Transactions on Computers*, *C-22*, 442–450.

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

Cattell, R. (1976). Constraints on movement rules. *Language*, *52*, 18–50.

Chien, Y.-C., & Wexler, K. (1990). Children's knowledge of locality conditions in binding as evidence for the modularity of syntax and pragmatics. *Language Acquisition*, *1*, 225–295.

Chomsky, N. (1973). Conditions on transformations. In S. Anderson & P. Kiparsky (Eds.), *A festschrift for Morris Halle* (pp. 232–286). New York: Holt, Rinehart, & Winston.

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

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

Christiansen, M. H., & MacDonald, M. C. (2009). A usage-based approach to recursion in sentence processing. *Language Learning*, *59*(Suppl. 1), 126–161.

Chung, S., Ladusaw, W., & McCloskey, J. (in press). Sluicing(:) between structure and inference. In R. G. Bravo, L. Mikkelsen, & E. Potsdam (Eds.), *Representing language: Essays in honor of Judith Aissen* (pp. 285–296). Santa Cruz, CA: LRC Publications.

Chung, S., & McCloskey, J. (1983). On the interpretation of certain island facts in GPSG. *Linguistic Inquiry*, *14*, 704–713.

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.

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, *386*, 604–608.

Conroy, A., Takahashi, E., Lidz, J., & Phillips, C. (2009). Equal treatment for all antecedents: How children succeed with Principle B. *Linguistic Inquiry*, *40*, 446–486.

Cowan, N. (2000). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–185.

Culicover, P., & Postal, P. (Eds.). (2001). *Parasitic gaps*. Cambridge, MA: MIT Press.

Deane, P. (1991). Limits to attention: A cognitive theory of island constraints. *Cognitive Linguistics* *2*, 1–63.

De Villiers, J., Cahillane, J., & Altreuter, E. (2006). What can production reveal about Principle B? In K. Ud Deen, J. Nomura, B. Schulz, & B. D. Schwartz (Eds.), *The proceedings of the inaugural conference on Generative Approaches to Language Acquisition – North America* (pp. 89–100). University of Connecticut Occasional Papers in Linguistics #4.

den Dikken, M. (2001). Pluringulars, pronouns, and quirky agreement. *The Linguistic Review*, *18*, 19–41.

Eberhard, K. M., Cutting, J. C., & Bock, J. K. (2005). Making syntax of sense: Number agreement in sentence production. *Psychological Review*, *112*, 531–559.

Elbourne, P. (2005). On the acquisition of Principle B. *Linguistic Inquiry*, *36*, 333–365.

Elbourne, P. (2008). The interpretation of pronouns. *Language and Linguistics Compass*, *2*, 119–150.

Engdahl, E. (1982). Restrictions on unbounded dependencies in Swedish. In E. Engdahl & E. Ejerhed (Eds.), *Readings on unbounded dependencies in Scandinavian languages* (pp. 151–174). Stockholm: Almqvist & Wiksell.

Engdahl, E. (1983). Parasitic gaps. *Linguistics and Philosophy*, *5*, 5–34.

Engdahl, E. (1997). Relative clause extractions in context. *Working Papers in Scandinavian Syntax*, *60*, 51–79.

Erteschik-Shir, N. (1973). *On the nature of island constraints* (PhD dissertation). Massachusetts Institute of Technology, Cambridge, MA.

Evans, G. (1980). Pronouns. *Linguistic Inquiry*, *11*, 337–362.

Fodor, J. D. (1978). Parsing strategies and constraints on transformations. *Linguistic Inquiry*, *9*, 427–473.

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

Franck, J., Lassi, G., Frauenfelder, U., & Rizzi, L. (2006). Agreement and movement: A syntactic analysis of attraction. *Cognition*, *101*, 173–216.

Frazier, L. (1985). Syntactic complexity. In D. Dowty, L. Karttunen, & A. Zwicky (Eds.), *Natural language processing: Psychological, computational, and theoretical perspectives* (pp. 129–189). Cambridge: Cambridge University Press.

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.

Gibson, E., & Thomas, J. (1999). Memory limitations and structural forgetting: The perception of complex ungrammatical sentences as grammatical. *Language and Cognitive Processes*, *14*, 225–248.

Goodluck, H., & Rochemont, M. S. (1992). *Island constraints: Theory, acquisition, and processing*. Dordrecht: Kluwer.

Gordon, P. C., & Hendrick, R. (1997). Intuitive knowledge of linguistic coreference. *Cognition*, *62*, 325–370.

Grimshaw, J., & Rosen, S. (1990). Knowledge and obedience: The developmental status of the binding theory. *Linguistic Inquiry*, *21*, 187–222.

Grodzinsky, Y., & Reinhart, T. (1993). The innateness of binding and coreference. *Linguistic Inquiry*, *24*, 69–101.

Grosu, A. (1973). On the status of the so-called right roof constraint. *Language*, *49*, 294–311.

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

Hagoort, P., Wassenaar, M., & Brown, C. (2003). Syntax-related ERP effects in Dutch. *Cognitive Brain Research*, *16*, 38–50.

Han, C., & Kim, J. (2004). Double relative clauses'' in Korean? *Linguistic Inquiry*, *35*, 315–337.

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.

Hartsuiker, R. J., Antón-Méndez, I., & Van Zee, M. (2001). Object attraction in subject-verb agreement construction. *Journal of Memory and Language*, *45*, 546–572.

Hartsuiker, R. J., Schriefers, H. J., Bock, J. K., & Kikstra, G. M. (2003). Morphophonological influences on the construction of subject-verb agreement. *Memory and Cognition*, *31*, 1316–1326.

Hawkins, J. A. (1999). Processing complexity and filler-gap dependencies across languages. *Language*, *75*, 224–285.

Heim, I. (1998). Anaphora and semantic interpretation: A reinterpretation of Reinhart's approach. In U. Sauerland & O. Percus (Eds.), *The interpretive tract* (pp. 205–246). Cambridge, MA: MIT Working Papers in Linguistics.

Hendriks, P., Banga, A., van Rij, J., Cannizzaro, G., & Hoeks, J. (2010). *Adults' on-line comprehension of object pronouns in discourse*. Manuscript submitted for publication.

Hendriks, P., & Spenader, J. (2005/2006). When production precedes comprehension: An optimization approach to the acquisition of pronouns. *Language Acquisition*, *13*, 319–348.

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

Hsu, C.-C. N. (2006). *Issues in head-final relative clauses in Chinese: Derivation, processing, and acquisition* (PhD dissertation). University of Delaware, Newark, DE.

Huang, C.-T. J. (1982). *Logical relations in Chinese and the theory of grammar* (PhD dissertation). Massachusetts Institute of Technology, Cambridge, MA.

Hyams, N., & Wexler, K. (1993). The grammatical basis of null subjects in child language. *Linguistic Inquiry*, *24*, 421–459.

Ishizuka, T. (2009). *CNPC violations and possessor raising in Japanese*. Los Angeles, CA, UCLA.

Jakubowicz, C. (1984). On markedness and binding principles. In C. Jones & P. Sells (Eds.), *Proceedings of NELS 14* (pp. 154–182). Amherst, MA: University of Massachusetts Graduate Linguistics Student Association.

Kaan, E. (2007). Event-related potentials and language processing: A brief overview. *Language and Linguistics Compass*, *1*, 571–591.

Kaan, E., Choi, J., Rowlings, B., Ballantyne, J., & Wijnen, F. (2009, March). *Cognitive control and agreement.* Poster presented at the 22nd annual CUNY Conference on Human Sentence Processing, Davis, CA.

Kaufman, D. (1988). *Grammatical and cognitive interactions in the study of children's knowledge of binding theory and reference relations* (PhD dissertation). Temple University, Philadelphia, PA.

Kayne, R. (1983). Connectedness. *Linguistic Inquiry*, *14*, 223–249.

Kayne, R. (1989). Notes on English agreement. *Central Institute of English and Foreign Languages Bulletin*, *1*, 41–67.

Kazanina, N. (2005). *The acquisition and processing of backwards anaphora* (PhD dissertation). University of Maryland, College Park, MD.

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

Kazanina, N., & Phillips, C. (2010). Differential effects of constraints in the processing of backwards anaphora. *Quarterly Journal of Experimental Psychology, 63*, 371–400.

Kempson, R., & Kiaer, J. (2010). Multiple long-distance scrambling: Syntax as reflections of processing. *Journal of Linguistics, 46*, 127–192.

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.

Kimball, J., & Aissen, J. (1971). I think, you think, he think. *Linguistic Inquiry, 2*, 241–246.

Kluender, R. (1998). On the distinction between strong and weak islands: A processing perspective. In P. Culicover & L. McNally (Eds.), *Syntax and semantics 29: The limits of syntax* (pp. 241–279). San Diego, CA: Academic Press.

Kluender, R. (2005). Are subject islands subject to a processing account? In V. Chand, A. Kelleher, A. J. Rodriguez, & B. Schmeiser (Eds.), *Proceedings of the 23rd West Coast conference on linguistics* (pp. 475–499). Somerville, MA: Cascadilla Press.

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

Krifka, M. (1990). Four thousand ships passed through the lock: Object-induced measure functions on events. *Linguistics and Philosophy, 13*, 487–520.

Kuno, S. (1973). *The structure of the Japanese language*. Cambridge, MA: MIT Press.

Kuno, S. (1976). Subject, theme, and speaker's empathy – a reexamination of relativization phenomena. In C. Li (Ed.), *Subject and topic* (pp. 417–444). New York: Academic Press.

Kush, D. (2010). *On the escapability of islands in Scandinavian*. Unpublished manuscript, University of Maryland, College Park, MD.

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

Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science, 29*, 1–45.

Lewis, R. L., Vasishth, S., & van Dyke, J. A. (2006). Computational principles of working memory in sentence comprehension. *Trends in Cognitive Sciences, 10*, 447–454.

McKee, C. (1992). A comparison of pronouns and anaphors in Italian and English acquisition. *Language Acquisition, 2*, 21–54.

Manning, C., & Schütze, H. (1999). *Foundations of statistical natural language processing*. Cambridge, MA: MIT Press.

Maratsos, M., & Kowalsky, M. (2005, January). *Mild subjacency violations and awkward grammatical sentences receive similar grammatical acceptability judgments from speakers*. Talk presented at the Linguistic Society of America annual meeting, Oakland, CA.

Matsuoka, K. (1997). Binding conditions in young children's grammar: Interpretations of pronouns inside conjoined NPs. *Language Acquisition, 6*, 37–48.

Mazuka, R., Jincho, N., & Oishi, H. (2009). Development of executive control and language processing. *Language and Linguistics Compass, 3*, 59–89.

McDaniel, D., Cairns, H., & Hsu, J. (1990). Binding principles in the grammars of young children. *Language Acquisition, 1*, 121–139.

McElree, B., & Griffith, T. (1998). Structural and lexical constraints on filling gaps during sentence comprehension: A time-course analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*, 432–460.

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

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.

Nordgaard, T. (1985). *Word order, binding, and the empty category principle* (Cand. Philol. Thesis). University of Trondheim, Trondheim, Norway.

Omaki, A., & Schulz, B. (in press). Filler-gap dependencies and island constraints in second language sentence processing. *Studies in Second Language Acquisition*.

Osterhout, L., & Mobley, L. A. (1995). Event-related brain potentials elicited by failure to agree. *Journal of Memory and Language*, 34, 739–773.

Phillips, C. (1995). Syntax at age two: Cross-linguistic differences. In C. Schütze, J. Ganger, & K. Broihier (Eds.), *Papers on Language Acquisition and Processing (MIT Working Papers in Linguistics, volume 26)* (pp. 225–282). Reprinted (2010) in Language Acquisition, 17, 70–120.

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

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

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

Phillips, C., & Wagers, M. (2007). Relating structure and time in linguistics and psycholinguistics. In G. Gaskell (Ed.), *Oxford handbook of psycholinguistics* (pp. 739–756). Oxford: Oxford University Press.

Poeppel, D., & Wexler, K. (1993). The full competence hypothesis of clause structure in early German. *Language*, 69, 1–33.

Poulton, E. C. (1954). Eye-hand span in simple serial tasks. *Journal of Experimental Psychology*, 47, 403–410.

Pritchett, B. L. (1991). Subjacency in a principle-based parser. In R. C. Berwick (Ed.), *Principle-based parsing: Computation and psycholinguistics* (pp. 301–345). Dordrecht, The Netherlands: Kluwer.

Reinhart, T. (1983). *Anaphora and semantic interpretation*. London: Croom Helm.

Reinhart, T. (2006). *Interface strategies: optimal and costly computations*. Cambridge, MA: MIT Press.

Reinhart, T., & Reuland, E. (1993). Reflexivity. *Linguistic Inquiry*, 24, 657–720.

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

Rizzi, L. (1994). Some notes on linguistic theory and language acquisition: The case of root infinitives. *Language Acquisition*, 3, 371–393.

Ross, J. R. (1967). *Constraints on variables in syntax* (PhD dissertation). Massachusetts Institute of Technology.

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

Sakai, H. (1994). Complex NP constraint and case conversions in Japanese. In M. Nakamura (Ed.), *Current topics in English and Japanese* (pp. 179–203). Tokyo: Hituzi Syobo.

Schlenker, P. (1999). La flexion de l'adjectif en allemand: La morphologie de haut en bas. *Recherches Linguistiques de Vincennes*, 28, 115–132.

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

Sohn, H. (1980). Theme prominence in Korean. *Korean Linguistics*, 2, 2–19.

Spenader, J., Smits, E.-J., & Hendriks, P. (2009). Coherent discourse solves the pronoun interpretation problem. *Journal of Child Language*, 36, 23–52.

Sprouse, J., Wagers, M., & Phillips, C. (2010). *A test of the relation between working memory capacity and syntactic island effects.* Manuscript submitted for publication.

Staub, A. (2009). On the interpretation of the number attraction effect: Response time evidence. *Journal of Memory and Language, 60*, 308–327.

Steiner, I. (2009). Partial agreement in German: A processing issue? In S. Featherston & S. Winkler (Eds.), *The fruits of empirical linguistics* (Vol. II, pp. 239–260). Berlin: de Gruyter.

Stevens, S. S. (1957). On the psychophysical law. *Psychological Review, 64*, 153–181.

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

Taraldsen, K. T. (1982). Extraction from relative clauses in Norwegian. In E. Engdahl & E. Ejerhed (Eds.), *Readings on unbounded dependencies in Scandinavian languages* (pp. 205–221). Stockholm: Almqvist & Wiksell.

Thornton, R., & Wexler, K. (1999). *Principle B., VP-ellipsis, and interpretation in child grammar.* Cambridge, MA: MIT Press.

Tokimoto, S. (2009). *Island phenomenon in Japanese and working memory: Syntactic constraints independent from working memory constraints.* Poster presented at the 22nd annual CUNY sentence processing conference, Davis, CA.

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.

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*, 97–106.

Trueswell, J., Sekerina, I., Hill, N. M., & Logrip, M. L. (1999). The kindergarten path effect: Studying on-line sentence processing in children. *Cognition, 73*, 89–134.

van Gompel, R., & Liversedge, S. (2003). The influence of morphological information on cataphoric pronoun assignment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 128–139.

Vigliocco, G., Butterworth, B., & Garrett, M. F. (1996). Subject-verb agreement in Spanish and English: Differences in the role of conceptual constraints. *Cognition, 61*, 261–298.

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.

Wellwood, A., Pancheva, R., Hacquard, V., Fults, S., & Phillips, C. (2009, March). *The role of event comparison in comparative illusions.* Poster presented at the CUNY conference on human sentence processing, Davis, CA.