

Processing Accounts for Superiority Effects

Ivan A. Sag
Stanford University

Philip Hofmeister
UC San Diego

Inbal Arnon Neal Snider
Stanford University

Florian Jaeger
University of Rochester

March 9, 2008

1 Introduction

At least since Kuno & Robinson 1972, it has been widely assumed within generative discussions that contrasts like (1a) vs. (1b) are to be explained by principles of grammar, indeed principles of *Universal Grammar*, such as Chomsky's (1973) 'Superiority' condition or the Minimal Link Condition of Chomsky (1995: 296):

- (1) a. Who __ broke what?
- b. *What did who break __ ?

However, in an important early paper, Bolinger (1978) argued that in the right context, examples like (2) are in fact grammatical, as are examples like the following:

- (2) I know what just about everybody was ASKED to do, but what did who (actually) DO __ ?

These Superiority 'violations' (henceforth SUVs) were subsequently explained away by Pesetsky (1987), along with examples like (3) [Karttunen (1977)], as the result of a distinct grammatical mechanism which he called 'D(iscourse)-Linking'.

- (3) Which newspaper did which student read __ ?

To this day, most generative discussions of *wh*-constructions¹ assume that Superiority Effects (SEs) are the norm, to be explained by principles of UG, and that SUVs are a secondary phenomenon, to be analyzed via D-Linking, or some other mechanism peripheral to the central concerns of grammatical theory.

In the alternative approach to SUVs considered here, examples like (1b) and (2) are all treated as fully *grammatical* and are analyzed uniformly by the grammar of *wh*-interrogatives, which includes neither a SE-inducing condition nor any D-Linking apparatus. The reduced *acceptability* of examples like (1b), examined more closely below, is explained in terms of the interaction of factors known on independent grounds to contribute to processing difficulty, which in turn leads to degraded acceptability.²

2 D-linking

To preserve the generalization behind Superiority, the notion of D-linking provides a crucial escape hatch to account for examples like (3). That is, the grammar retains the relevant constraints on *wh*-phrase order and examples like (1b) are ruled out as grammatically ill-formed. A formal notion of D-linking, however, is never provided by Pesetsky, nor does one appear in the secondary literature. Perhaps the most direct indication of its meaning appears in the following text from Pesetsky (2000):

“Context sets previously mentioned in the discourse qualify a phrase as D-linked, but so do sets that are merely salient (e.g. *which book*, in a context where speaker and hearer both know that reference is being made to a reading list for a course) and sets whose salience is culturally determined (e.g. *what day of the week, which sign of the zodiac*)” [p. 16].

According to this passage, individual lexical items and phrases are not inherently specified for D-linking. Instead, this property emerges from the context or conventionalized interpretations, and hence nothing theoretically prevents bare *wh*-items like *who* and *what* from being D-linked. Note that it is hard to reconcile Pesetsky’s proposal as stated here, where the grammaticality of examples like (1b) is claimed to vary from context to context, with the practice of the syntactic literature, where such SUVs are treated as ungrammatical, *tout court*. Moreover, the above quote offers little clarity on the perceived contrast between examples like (3) and (1b), since both examples have no prior context whatsoever and the range of books in (3) is just as unknown to the reader as the objects that *what* ranges over in (1b).

In addition, within the theory of grammar assumed by Pesetsky, it is difficult to see why contextualization should be able to invalidate an otherwise universal constraint such as Superiority. Pesetsky acknowledges, in fact, the absence of any rationale for why D-linking should create an exception to the constraints on *wh*-phrase order:

“A reliable rule of thumb is that if a *wh*-word in a multiple question can be felicitously paraphrased with an expression of the form *which of the X*, it can cause the Superiority effect to disappear. The reason for this link between semantics and syntax is obscure, and will remain obscure even at the end of this book” [p. 16].

This concern is in fact anticipated by Chung (1994:39) who asks, “[W]hat is it about the ability to narrow down the domain of *wh*-quantification ‘enough’ that makes it possible for strict locality to be violated?” Since there is no explanation for why contextualization voids this one syntactic constraint, D-linking ultimately stands as a mere descriptive label for examples which run counter to the predictions of Superiority.

Bolinger (1978) also discusses contrasts like (4b,c), which have proven to be problematic for standard syntactic accounts of SEs:

- (4) a. Who took what where?
b. What did who take ___ where?

- c. Where did who take what ___ ?

To analyze examples like these, Pesetsky (2000: 48ff) assumes that a *wh*-expression can be marked as ‘invisible’, which allows it to undergo ‘*wh*-feature movement’, rather than overt *wh*-movement. Pesetsky proposes to account for (4b,c) in terms of Richards’ (1997) Principle of Minimal Compliance. The nonsubject *wh*-expressions in these examples may undergo overt movement, because the prior covert movement of the subject suffices for compliance with the Superiority Condition (the ‘Attract Closest’ Condition). Quite apart from the theoretical difficulties raised by any principle such as Richards’, this account is insufficiently general, covering none of the preceding SUVs involving only two *wh*-expressions. It also fails to account for well-formed SUV examples like the following, discussed by Kennedy (2005):

- (5) Who knows what who bought?

The account in terms of the Principle of Minimal Compliance is both theoretically implausible and lacking in independent motivation. Hence, as Pesetsky himself admits (2000: 104), his account stands as a ‘mere placeholder for a future explanation of the phenomenon’.

Finally, some predictions that the theory of D-linking makes seem to be empirically wrong, as we demonstrate in the next section. SUVs occur in naturally spoken and written data where both *wh*-phrase arguments are bare, where no prior context set has been mentioned, and where no answer set has previously been made salient. On the basis of these points, and the absence of any other grammatical explanation, there remains no principled, grammar-based account of the contrast between (3) and examples like (1b). This leads us to entertain the view that SUVs like (1b) may, in fact, be entirely grammatical (i.e. in accordance with all principles of grammar), but rare and of reduced acceptability for reasons that are independent of syntax.

3 SUVs in the Wild

Chief among the reasons for revisiting the status of SUVs is the fact that the empirical data conflicts with the claims of previous analyses. The Superiority condition (or a more abstract principle that derives the same effects (like the MLC) predicts the general absence of naturally-occurring SUVs or perhaps that any such occurrences will have the status of performance-related errors. As noted by Arnon et al. (2005) and by Clifton et al. (2006), however, SUVs (even with two bare *wh*-phrases) can easily be found on the internet via Google searches.³ Of course, internet data is noisy for many reasons, and one cannot conclude that sentences are grammatical simply because one can find them on the internet. But the nature of these examples and the contexts in which such data occur are quite revealing. First of all, attested examples appear in written and presumably edited text, a fact that speaks against their status as errors:

- (6) a. What did who know and when did they know it?
[<http://www.antigonishreview.com/bi-113/113-curb.html>]

- b. What did who say and who did the asserting?
[<http://www.thenation.com/doc/20030512/cockburn>]

Consider, as well, the following examples:⁴

- (7) a. Although nothing on this planet (or any other) can compete with the utter horror that is cilantro! Where the heck did who the heck come up with adding that gawdawful weed to otherwise civilized hote-cue-zeen?
[<http://www.scrappleface.com/MT/archives/001655.html>]
b. Why the hell did who change the good, suitable icon it had (like on v 1.74) to that stupid eGarbage!?
[<http://www.versiontracker.com/dyn/moreinfo/macosx/10658&page=2>]

Pesetsky (1987) coins the term ‘aggressively non-D-Linked’ for cases like *who the heck*, which presumably entails that attested examples like (7a) are inconsistent with the intended notion of D-Linking. Similarly, given the fact that the recriminations were just starting in the following example, it is unlikely that the answer set associated with *what* is more salient than the one associated with *who*:

- (8) We watch for a few more minutes as the recriminations begin: what did whom [*sic*] say and what did who hear? [*Tomato is Coming*, by Crispin Oduobuk]

Other examples are similar – there is no evidence that the prior discourse narrows down the range of possible answers for the higher *wh*-expression:

- (9) What, do you think this is a game? What rules should who follow? This shit-sandwich is a reality – a competition for survival between all souls...
[<http://www.cruel.com/discuss/viewTopic.php/83308>]
(10) I must have missed something. What did who do to Pierre Salinger?
[<http://www.freerepublic.com/forum/a3b1c8a4d1847.htm>]

What the fronted *wh*-expressions in these sentences have in common, however, is the fact that the initial *wh*-expression continues the topic thread and generally asks the more pertinent and relevant question in relation to the previous discourse.^{5,6} Once one sees this, it is not hard to construct contextualized examples where an SUV is more acceptable than the corresponding example that obeys Superiority. For instance, (11a) is plainly more acceptable than (11b):

- (11) a. You’re a complete mess... What did who DO to you when you were a child?!
b. You’re a complete mess... Who did what to you when you were a child?!

SUVs of the sort documented here, of course, are relatively rare when compared to multiple *wh*-questions which do not violate Superiority or where one or more *wh*-phrase is not a bare *wh*-word. Despite these differences in frequency or usage, we do find numerous examples of multiple *wh*-questions that violate Superiority and which appear to occur in spite of the absence of discourse properties associated with D-linking. Accordingly, the empirical data favors treating SUVs of all sorts as fully grammatical, and thus compels us to search outside of grammar for an explanation of the differences in frequency (and acceptability).

4 Acceptability

In addition to the occurrence of SUVs in written and spoken language, another motivation for reconsidering the treatment of SUVs concerns the systematic variation in acceptability judgments of SUVs. The original version of the D-linking theory presented in Pesetsky (1987) claims that it is the in-situ subject *wh*-phrase that must be D-linked to avoid an SE. Subsequently, Pesetsky (2000: 97) revises this opinion, suggesting that “either the higher *wh*-phrase or the *wh*-phrase that moves overtly must be D-linked,” although no rationale is presented for why either *wh*-phrase can rescue an SUV. This raises the question of how acceptability relates to the content of the *wh*-expressions in the multiple argument positions of SUVs.

We investigated this issue using Magnitude Estimation (ME), a technique allowing subjects to construct a structured space of systematic judgments, where a given stimulus is evaluated with respect to a chosen norm. ME studies have been demonstrated [see Bard et al. 1996 and Keller 2000, *inter alia*] to be reliable for studying sentence acceptability judgments. In a series of controlled acceptability studies using ME,⁷ we manipulated whether the *wh*-phrases appearing in embedded SUVs were bare *wh*-phrases or “complex” multi-word *wh*-phrases – *which*- \bar{N} phrases, but also *what*- \bar{N} phrases. The acceptability results we obtained verify that a complex *wh*-phrase in either argument position of an SUV significantly improved acceptability. That is, a single complex *wh*-phrase in the object position or subject position of an SUV raised acceptability judgments significantly, as compared to an SUV with zero complex *wh*-phrases. Figure 1 illustrates that both *what*- \bar{N} phrases and *which*- \bar{N} phrases in object position improve acceptability, when contrasted with a bare *wh*-word. Note that this shows again that it is not only *which*-phrases that ameliorate the degradedness of SUVs; rather, the results suggest that the complexity of the *wh*-phrase is a critical factor affecting acceptability judgments. Of course, all examples were judged without any preceding context, meaning that all *wh*-expressions were equally D-linked in terms of contextual salience. Also, note that having a complex *wh*-phrase in either position (subject or object) raised acceptability levels.

Crucially, SUVs with two complex *wh*-phrases (e.g. *which book did which student read*) were judged significantly better than SUVs with only one complex *wh*-phrase. This additivity goes unexplained in analyses that appeal to exception principles like D-linking. As we shall see, however, there is a clear and principled reason for expecting such additivity on the hypothesis that the degraded acceptability of SUVs results from cumulative processing pressures.

Hofmeister et al. (2007) additionally investigated the extent to which processing measures correlate with acceptability. Using the self-paced reading methodology, the effect of *wh*-phrase complexity on processing was examined by manipulating the kind of *wh*-expression appearing in each argument position. As with the acceptability task, crossing the complexity of the object *wh*-expression with the complexity of the in-situ subject phrase produced four conditions per experimental item. A sample item is shown in (12):

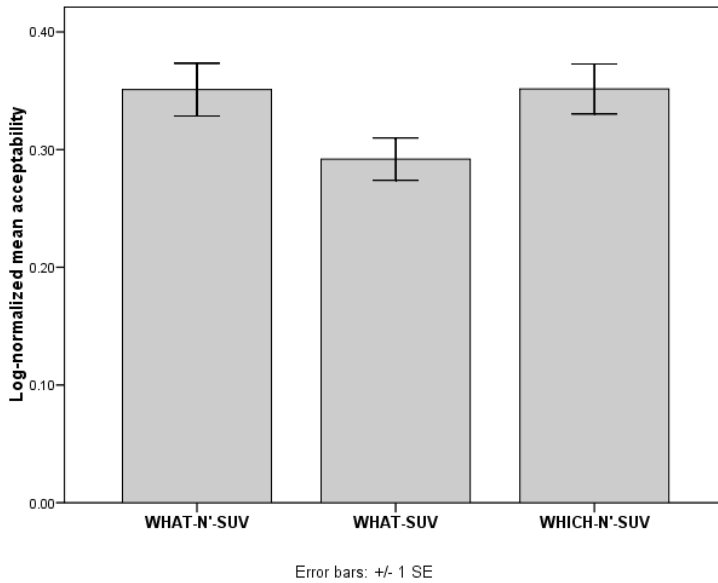


Figure 1: Log-normalized acceptability ratings of three kinds of Superiority-violations, varying by type of fronted *wh*-phrase, reported in Hofmeister et al. 2007

- (12)
- a. Ashley disclosed what who signed after receiving permission from the president.
 - b. Ashley disclosed which agreement who signed after receiving permission from the president.
 - c. Ashley disclosed what which diplomat signed after receiving permission from the president.
 - d. Ashley disclosed which agreement which diplomat signed after receiving permission from the president.

Paralleling the acceptability judgments, less complex (and less informative) *wh*-phrases in both object and subject position resulted in slower reading at the verb and spillover region, as shown in Figure 2. On average, reading times at the verb for examples like (12a) were slowest, while sentences like (12d) produced the fastest times. The other two conditions, each with one bare *wh*-word and one *which-N* phrase, were intermediate between these two extremes. In other words, the reading time results show main effects of the kind of *wh*-phrase used in both positions of the SUVs.

In essence, the results of the acceptability tasks and the comprehension-based experiment align with one another. The condition with the fastest processing times also received the highest mean rating of acceptability. Conversely, the hardest condition for processing was judged the worst in terms of acceptability. Again, the intermediate conditions in terms of processing yielded similarly intermediate scores in terms of acceptability. These strong parallels between processing and acceptability support an account that attributes the relative unacceptability of SUVs (and non-SUVs) to the degree of




Figure 2: Mean residual reading times at verb plus spillover region in SUVs, based on results reported in Hofmeister et al. 2007

processing difficulty. Grammar-based accounts, by contrast, fail to predict the observed parallels.

5 Processing Factors

Based on the foregoing reasons, there is every reason to look outside of grammar for an account of whatever deviance inheres in the SUV examples routinely cited in the syntactic literature. Generally speaking, unacceptability is understood as resulting from both properties of the competence grammar and performance-related processing factors. In a well-known consideration of the acceptability effects of the latter, Chomsky & Miller (1963) pursued an explanation of the unacceptability of complex center-embeddings (e.g. *The boy the girl the host knew brought left*) based on cognitive limitations. According to such an analysis, the grammar effectively overgenerates the set of possible embeddings, but the members of that set that are found acceptable are limited by constraints on cognitive resources.

A similar type of reasoning can be applied to the case of SUVs. On such an account, processing pressures rise past some threshold during the processing of SUVs, such that the token is considered unacceptable. A *which-N* phrase or some equally complex *wh*-phrase then must partly ameliorate this processing difficulty. This story then raises the following questions: (1) why should SUVs be inherently harder to process than non-SUVs and (2) why should a complex *wh*-phrase facilitate processing?

Answers to both of these questions are evident in the psycholinguistic literature on dependency processing. Despite the fact that the distance between filler and gap in *wh*-

constructions is, theoretically speaking, unbounded, psycholinguistic research provides strong evidence that processing difficulty increases with dependency length (Frazier & Fodor 1978, King & Just 1991, Gibson 1998, 2000, Hawkins 2004, Grodner & Gibson 2005, Fiebach et al. 2001). In the relative clause constructions below, for instance, the increased distance between the subject NP and the gap site causes increased processing difficulty around this retrieval point, as determined by various behavioral measures:

- (13) a. The administrator who_i the nurse supervised -- _i scolded the medic.
 b. The administrator who_i the nurse who was from the clinic supervised -- _i scolded the medic.

In non-SUV word orders, the object *wh*-expression is immediately contiguous to its verbal head. But SUV *wh*-orders create a longer dependency between the object phrase and its subcategorizing head. Thus, distance-based theories of sentence processing predict higher processing costs for the SUV word order.

By itself, this increased dependency length would be unlikely to cause severe processing issues, as many constructions with far longer dependencies are perfectly acceptable in English. However, there are other pressures associated with SUVs that quite plausibly increase their difficulty. First, processing a question requires the computation of existential presuppositions, i.e. the presupposition that there is an answer of the appropriate sort (Kroch 1989). Hence, a multiple *wh*-question involves the computation of two such presuppositions, or a single, doubly quantified presupposition. Arnon et al. (2005) point out that unary *wh*-questions and non-SUV multiple *wh*-questions receive much lower acceptability ratings than declaratives without any internal filler-gap dependencies. Since there is no debate about the grammatical status of these constructions, the relatively low acceptability must derive from a grammar-external source. That is, *wh*-dependencies, in general, and multiple *wh*-questions more specifically, impose significant processing costs, despite their unquestioned grammatical status.

SUVs add to these already significant costs by imposing the need to create a dependency inside another dependency, albeit for a brief time. In this sense, they are very much akin to center-embeddings, whose difficulty is tied to the challenges of maintaining multiple dependencies simultaneously. Such constructions require the accurate retention of multiple filler-phrases and the retrieval of the right phrase at the right time.

Moreover, it is worth noting the similarity of the two argument phrases (*what* and *who*) in multiple *wh*-questions. In (1b), for example, both *wh*-expressions are single *wh*-words that only differ semantically in terms of their animacy specifications; they share some obvious phonological characteristics, as well. A wealth of research in language processing has revealed that the retrievability of a mental representation varies with the presence of other representations in working memory that overlap the feature specifications of the target (Lewis 1996, Gordon et al. 2001, Gordon et al. 2002, Van Dyke & Lewis 2003, Gordon et al. 2006, inter alia). In particular, demands on retrieval are known to increase when a target shares features with similar representations. For instance, Gordon et al. (2006) demonstrate that object clefts are harder to read when the clefted NP and subject NP are both definites or both proper names, as in (15):

- (14) a. It was the barber that the lawyer saw in the parking lot.

- b. It was John that Bill saw in the parking lot.
- (15) a. It was the barber that Bill saw in the parking lot.
- b. It was John that the lawyer saw in the parking lot.

Notice that these constructions are similar to the configuration of arguments in SUVs with the object preposed in front of the subject NP. The similarity of the NPs in (14) impedes the retrieval process at the embedded verb, and so a similar reasoning can be extended to the case of SUVs like (1b): the similarity of the two, unintegrated *wh*-items creates an interference effect at the verb that slows the retrieval process. Incorporating similarity-based interference effects into an understanding of the acceptability of SUVs further accounts for the seriously degraded nature of sentences with two identical *wh*-expressions like *I know who who killed*.

Cumulatively, the presence of these pressures can explain why SUVs occur far less frequently than their non-SUV counterparts and why they are found to be relatively low in acceptability. Similarly, the improved acceptability of SUVs with *which*-phrases or other complex *wh*-expressions can be understood by appealing to other cognitive principles. In particular, Hofmeister (2007) demonstrates that the processing facilitation and greater acceptability associated with *which*- \bar{N} phrases is not specific to SUVs. A range of dependency environments in English display a sensitivity to the content of the filler phrases, including *wh*-islands, complex noun phrase constructions, subject condition violations and other island contexts. These effects, however, are not limited to islands: even in unquestionably grammatical contexts, *which*- \bar{N} phrases improve processing at retrieval (gap) sites, as compared to bare *wh*-words. Indeed, the empirical evidence suggests that the relevant property is not even specific to *wh*-dependencies. In dependencies with referential filler-phrases, more complex or informative filler-phrases facilitate processing at the retrieval site. For instance, definite NPs like the one in (16b), which is syntactically and semantically more informative than the definite in (16a), were found to produce faster reading times beginning at the subcategorizing verb (*encouraged*) (Hofmeister 2007):

- (16) a. The diplomat contacted **the dictator** who the activist looking for more contributions **encouraged** to preserve natural habitats and resources.
- b. The diplomat contacted **the ruthless military dictator** who the activist looking for more contributions **encouraged** to preserve natural habitats and resources.

In other words, the informativity of filler-phrases attenuates processing difficulty at the subcategorizing head. This appears to constitute a rather general property of language processing, as it occurs in a wide variety of contexts with a variety of phrase-types. To account for this pattern, Hofmeister (2007) argues that linguistic expressions that contain more syntactic and semantic information facilitate the retrieval from memory of the corresponding mental representations. This relationship is argued to exist because processing, categorizing, and integrating new information about an entity increases the associated activation level (Lewis & Vasisht 2005). Basically, thinking more about something makes it easier to remember. Additionally, the *result* of a building a more

elaborate mental representation is that it is more likely to be distinct from other representations in memory, and thus suffer less *interference*.⁸

By this logic, the facilitated processing in SUVs with *which-N* phrases occurs because they elicit more syntactic and semantic processing than bare *wh*-words. Hence, the corresponding mental representations are more activated and thus easier to retrieve at the verb. If the difficulty of processing a lexical head is a function of (among other things) the activation levels of its arguments, then the form preferences *wh*-questions emerge as a preference for high argument activation at the point when the head whose argument is ‘extracted’ is being processed. A preference for high activation for all arguments consequently explains the additive effects of *which-N* phrases in SUVs. This perspective also agrees with the study of Frazier & Clifton (2002), who show that *which-N* phrases are better antecedents for pronouns than bare *wh*-words like *who* and *what*. Since highly activated representations encourage the use of “high accessibility anaphors” like pronouns (Ariel 1990), this supports the conclusion that processing a *which-N* phrase leads to a higher level of argument activation than processing a bare *wh*-phrase.

6 Conclusion

Research on language processing contains a number of principled reasons for why we should expect SUVs to engender serious processing difficulty. Each of these processing burdens alone does not constitute an insurmountable challenge, but when they all coincide, they pose significant problems for rapid, on-line processing. Importantly, none of these processing explanations are specific to SUVs: they represent general and independently motivated processing principles that affect language comprehension because of limitations on cognitive resources. In this sense, treating the low acceptability of SUVs as the result of processing difficulties avoids introducing any new explanatory machinery.

More generally, this manner of understanding variation in acceptability has the advantage of uniting the analyses of variation in a range of linguistic constructions, as is done in much recent work (see, e.g. Hawkins 2004, Race & MacDonald 2003, and Ferreira & Dell 2000). We likened SUVs to center-embeddings because they share a number of key features, including the fact that both constructions exhibit variation in acceptability as a function of non-structural manipulations. Such analyses can then be tested on similar cases where judgments of acceptability vary substantially. By contrast, mechanisms like D-linking make highly construction-specific statements and hence are inherently unable to explain acceptability differences in other constructions. Quite apart from its empirical deficiencies, the ad hoc nature of D-Linking makes it a priori less desirable as a hypothesis.

A processing-based explanation of SUVs also naturally captures variation among speakers and even for the same speaker over time. Individual differences in working memory capacity have the potential to account for the fact that some individuals are much more accepting of SUVs than others. Besides differences *across* individuals, the incumbent processing pressures on a single language user also vary widely throughout discourse or reading. Other on-going cognitive tasks, fatigue, or even just absent-mindedness may limit an individual’s ability to process an incoming string. Moreover,

various contextual manipulations may indeed make processing an SUV easier, as suggested by the D-linking hypothesis, including previous mentions of the possible answer set. These facts, too, follow directly from a performance-based story of these effects, as repeated mentions boost the activation level of a given entity. Hence, our conclusions can be seen as actually corroborating the basic intuitions behind D-linking, but they provide a more detailed, functionally-driven account of those intuitions, while expanding upon the set of factors that influence acceptability.

Processing-based explanations of acceptability, however, must confront certain challenges. For instance, the present account must confront the non-universality of SEs. Some languages (e.g. German) exhibit limited SEs, as documented in detail by Featherston (2005), while other languages (e.g. Russian) do not evidence any effects of Superiority at all (See Fedorenko & Gibson in preparation). Nevertheless, such differences make perfect sense under a processing account of SUVs, because of the unique cues each language provides for processing. Arnon et al. (in press), building on the Competition Model (developed in MacWhinney & Bates 1989, MacWhinney 2004 and elsewhere), offer an account of these differences in terms of the availability and reliability of case marking as a processing cue. English speaker/hearers have no such morphological cues, while German exhibits a more reliable and expansive set of cues for determining thematic assignments, and Russian even more so. Thus, as Arnon et al. (in press) argue, a processing-based account of SEs actually provides a more satisfying and explanatory account of the observed cross-linguistic differences than an account based on stipulated variations of grammar.

Our main goal here has been to show that the available empirical evidence lines up against an interpretation of SEs in terms of grammar-based constraints. Counter to the predictions of such theories, SUVs of all sorts show up in natural language settings. Furthermore, the systematic variation in acceptability and processing difficulty associated with these constructions has no adequate, coherent explanation within a strictly competence-based approach. Most importantly, independently-motivated principles of processing predict not only the general processing difficulty and low acceptability of SUVs, but also the variation tied to the complexity of the *wh*-expressions, contextualization, and speaker differences. Such an approach, in its appeal to established principles of cognition, holds the potential to provide an increased understanding of other cases of variation in acceptability, as well. To the extent that such explanations are successful, the burden on the competence grammar is accordingly lightened, opening the door for a more parsimonious, ‘minimal’ approach to grammar.

Notes

¹Notable exceptions are Chomsky (1995: 387, fn. 69) and Ginzburg & Sag (2000: 247ff.).

²Here we follow Chomsky & Miller (1963) and Fanselow & Frisch (2005). See also the empirical arguments made by Kluender (1992, 1998, 2004), Featherston (2005) and Keller (2000).

³In all such investigations, it is essential to set aside two well-known classes of freely occurring SUVs, namely reprise questions of the sort illustrated in (17):

- (17) a. [A: What did Przemysław see?]
What did WHO see? [Echo Reprise (rising intonation)]
b. [A: What did he see?]
What did WHO see? [Reference Reprise (falling intonation)]

For fuller discussion of such cases, see Comorovski (1996, Ch. 3) and especially Ginzburg & Sag (2000, Ch. 7).

⁴The results shown here were all obtained between November 2004 and March 2005.

⁵On the topichood of *wh*-phrases, see, e.g. Grohmann 1998, to appear for German and Jaeger 2004 for Bulgarian. A closely related proposal is Kuno & Takami's (1993:112) Sorting Key Hypothesis, according to which the left-most word in a *wh*-question represents the key for sorting relevant pieces of information in the answer.

⁶In cases of more than two *wh*-expressions, such as (8), the most relevant *wh*-expression still generally comes first, though the ordering of the subsequent *wh*-expressions appears less constrained. (For similar Slavic data, see Grewendorf 2001 and Bošković 2002, among others.)

⁷See Hofmeister et al. 2007 and Arnon et al. in press. These results were first presented in Arnon et al. 2005.

⁸On this line of reasoning, one expects some degree of similarity-based interference in SUVs with two *which-N* phrases. Indeed, this is what one finds when the head nouns are identical, as in (i):

- (i) *Which robot did which robot kill?*

But in most multiple *wh*-questions the two head nouns are distinct. This fact, together with the activation boosts provided by these more complex *wh*-expressions, more than compensates for any similarity-based interference effects.

References

- Ariel, Mira. 1990. *Accessing Noun Phrase Antecedents*. London: Routledge.
- Arnon, Inbal, Bruno Estigarribia, Philip Hofmeister, T. Florian Jaeger, Jeanette Pettibone, Ivan A. Sag and Neal Snider. 2005. Rethinking Superiority effects: a processing model. Poster presented at the CUNY Sentence Processing Conference, University of Arizona.
- Arnon, Inbal, Neal Snider, Philip Hofmeister, T. Florian Jaeger, and Ivan A. Sag. in press. Cross-linguistic variation in a processing account: the case of multiple *wh*-questions. *Proceedings of BLS 32*. Berkeley: Berkeley Linguistic Society.
- Bard, Ellen G., Dan Robertson, and Antonella Sorace. 1996. Magnitude estimation of linguistic acceptability. *Language*, 72: 32-68.
- Bolinger, Dwight D. 1978. Asking more than one thing at a time. In: H. Hiz (ed.), *Questions*. Dordrecht: D. Reidel. Pp. 107-150.

- Bošković, Željko. 2002. On multiple *wh*-fronting. *Linguistic Inquiry*, 33.3: 351–383.
- Chomsky, Noam, and George A. Miller. 1963. Introduction to the formal analysis of natural languages. In Luce, R. D., R. R. Bush and E. Galanter (eds.), *Handbook of Mathematical Psychology*, Vol. 2. New York: John Wiley and Sons. Pp. 269–321.
- Chomsky, Noam. 1973. Conditions on transformations. In Stephen Anderson and Paul Kiparsky (eds.), *A Festschrift for Morris Halle*. New York: Holt, Rinehart and Winston. Pp. 232–286.
- Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.
- Clifton, Charles, Gisbert Fanselow, and Lyn Frazier. 2006. Amnestying superiority violations: processing multiple questions. *Linguistic Inquiry*, 37.1: 51–68.
- Comorovski, Ileana. 1996. *Interrogative Phrases and the Syntax-Semantics Interface*, vol. 59 of *Studies in Linguistics and Philosophy*. Dordrecht: Kluwer.
- Fanselow, Gisbert, and Stefan Frisch. 2006. Effects of processing difficulty on judgments of acceptability. In Gisbert Fanselow, Caroline Féry, Matthias Schlesewsky, and Ralf Vogel (eds.), *Gradience in Grammar*. Oxford: Oxford University Press.
- Featherston, Sam. 2005. Magnitude estimation and what it can do for your syntax: Some *wh*-constraints in German. *Lingua*, 115.11: 1525–1550.
- Fedorenko, Evelina, M. Babyonyshev, and Edward Gibson. 2004. The nature of case interference in on-line sentence processing in Russian. In Keir Moulton and Matthew Wolf (eds.), *Proceedings of NELS 34*. Amherst MA: GLSA Publications.
- Fedorenko, Evelina, and Edward Gibson. in preparation. Syntactic parallelism as an account of cross-linguistic superiority effects. Manuscript, MIT.
- Ferreira, Victor S., and Gary S. Dell. 2000. The effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive Psychology*, 40: 296–340.
- Fiebach, Christian, Matthias Schlesewsky, and Angela Friederici. 2001. Syntactic working memory and the establishment of filler-gap dependencies: Insights from ERPs and fMRI. *Journal of Psycholinguistic Research*, 30.3: 321–338.
- Frazier, Lyn, and Charles Clifton. 2002. Processing d-linked phrases. *Journal of Psycholinguistic Research*, 31.6: 633–659.
- Frazier, Lyn, and Janet Dean Fodor. 1978. The sausage machine: A new two-stage parsing model. *Cognition*, 6.4: 291–325.
- Gibson, Edward. 1998. Linguistic complexity: locality of syntactic dependencies. *Cognition*, 68.1: 1–76.

- Gibson, Edward. 2000. The Dependency Locality Theory: a distance-based theory of linguistic complexity. In A. Marantz, Y. Miyashita, and W. O’Neil (eds.), *Image, Language, Brain: Papers from the First Mind Articulation Project Symposium*. Cambridge: MIT Press. Pp. 95–126.
- Ginzburg, Jonathan, and Ivan A. Sag. 2000. Interrogative Investigations: The Form, Meaning, and Use of English Interrogatives (CSLI Lecture Notes 123). Stanford, CA: CSLI Publications. [Distributed by U. Chicago Press]
- Gordon, Peter, Randall Hendrick, and Marcus Johnson. 2001. Memory interference during language processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27.6: 1411-23.
- Gordon, Peter, Randall Hendrick, Marcus Johnson, and Yoonhyoung Lee. 2006. Similarity-based interference during language comprehension: evidence from eye tracking during reading. *Journal of Experimental Psychology: Learning, Memory, Cognition*, 32.6: 1304–1321.
- Gordon, Peter, Randall Hendrick, and William Levine. 2002. Memory-load interference in syntactic processing. *Psychological Science*, 13: 425–430.
- Grewendorf, Günther. 2001. Multiple *wh*-fronting. *Linguistic Inquiry*, 32.1: 87–122.
- Grodner, Daniel, and Edward Gibson. 2005. Consequences of the serial nature of linguistic input for sentential complexity. *Cognitive Science*, 29: 261–290.
- Grohmann, Kleanthes K. 1998. Syntactic inquiries into discourse restrictions on multiple interrogatives. In *Groninger Arbeiten zur germanistischen Linguistik*, 42: 1–60.
- Grohmann, Kleanthes K. To appear. Top issues in questions: topics – topicalization – topicalizability. In Lisa L.-S. Cheng and Norbert Corver (eds.), *Wh-Movement Moving On*. Cambridge, MA: MIT Press.
- Hawkins, John. 2004. *Efficiency and Complexity in Grammars*. Oxford: Oxford University Press.
- Hofmeister, Philip. 2007. *Representational Complexity and Memory Retrieval in Language Comprehension*. Ph.D. Thesis, Stanford University.
- Hofmeister, Philip, T. Florian Jaeger, Ivan A. Sag, Inbal Arnon, and Neal Snider. 2007. Locality and accessibility in *wh*-questions. In Sam Featherston and Wolfgang Sternefeld (eds.), *Roots: Linguistics in Search of its Evidential Base*. Berlin: Mouton de Gruyter.
- Jaeger, T. Florian. 2004. Topicality and superiority in Bulgarian *wh*-questions. In Olga Arnaudova, Wayles Browne, Maria Luisa Rivero, and Danijela Stojanovic (eds.), *Proceedings of FASL-12: The Ottawa Meeting 2003*. Michigan Slavic Publications. Pp. 207–228.

- Karttunen, Lauri. 1977. Syntax and semantics of questions. *Linguistics and Philosophy*, 1:3–44.
- Keller, Frank. 2000. *Gradience in Grammar: Experimental and Computational Aspects of Degrees of Grammaticality*. Ph.D. Thesis, University of Edinburgh.
- Kennedy, Becky. 2005. Interpretive effects in multiple interrogation. *Journal of Pragmatics*, 37.1: 1–42.
- King, Jonathan, and Marcel A. Just. 1991. Individual differences in syntactic processing: The role of working memory. *Journal of Memory and Language*, 30.5: 580–602.
- Kluender, Robert. 1992. Deriving island constraints from principles of predication. In Helen Goodluck and Michael Rochemont (eds.), *Island Constraints: Theory, Acquisition and Processing*. Dordrecht: Kluwer. Pp. 223–258.
- Kluender, Robert. 1998. On the distinction between strong and weak islands: a processing perspective. In Peter Culicover and Louise McNally (eds.), *Syntax and Semantics Volume 29: The Limits of Syntax*. Elsevier. Pp. 241–279.
- Kluender, Robert. 2004. Are subject islands subject to a processing account? In Vineeta Chand, Ann Kelleher, Angelo J. Rodríguez, and Benjamin Schmeiser (eds.), *WCCFL 23: Proceedings of the 23rd West Coast Conference on Formal Linguistics*. Somerville, MA: Cascadilla Press. Pp. 475–499.
- Kroch, Anthony. 1989. Amount quantification, referentiality, and long *wh*-movement. Manuscript, University of Pennsylvania. [Appears in Dimitriadis et al. (eds.), 1998, *Penn Working Papers in Linguistics*”]
- Kuno, Susumu, and Jane Robinson. 1972. Multiple *wh*-questions. *Linguistic Inquiry*, 3: 463–487.
- Kuno, Susumu, and Ken-Ichi Takami. 1993. *Grammar and Discourse Principles: Functional Syntax and GB Theory*. University of Chicago Press.
- Lewis, Richard. 1996. Interference in short-term memory: The magical number two (or three) in sentence processing. *Journal of Psycholinguistic Research*, 25.1: 93–115.
- Lewis, Richard and Shravan Vasishth. 2005. An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, 29: 1–45.
- MacWhinney, Brian. 2004. A Unified model of language acquisition. In J. Kroll and A. De Groot (eds.), *Handbook of Bilingualism: Psycholinguistic Approaches*. Oxford: Oxford University Press.
- MacWhinney, Brian, and Elizabeth Bates. 1989. *The Crosslinguistic Study of Sentence Processing*. New York: Cambridge University Press.

- Pesetsky, David. 1987. *Wh-in-situ: movement and unselective binding*. In E. Reuland and A. ter Meulen (eds.), *The Representation of (In)Definiteness*. Cambridge, MA: MIT Press. Pp. 98–129.
- Pesetsky, David. 2000. *Phrasal Movement and Its Kin*. Cambridge, MA: MIT Press.
- Race, David S., and Maryellen C. MacDonald. 2003. The use of *that* in the production and comprehension of object relative clauses. *Proceedings of the 25th Annual Meeting of the Cognitive Science Society*.
- Richards, Norvin. 1997. *What Moves Where When in Which Language?*. Ph.D. Thesis, MIT. [Available from *MIT Working Papers in Linguistics*].
- Rudin, Catherine. 1988. On multiple questions and multiple *wh*-fronting. *Natural Language and Linguistic Theory*, 6: 445–502.
- Van Dyke, Julie, and McElree, Brian. 2006. Retrieval interference in sentence comprehension. *Journal of Memory and Language*, 55: 157-166.