Syntactic Flexibility in the Noun: Evidence from Picture Naming

Nicholas A. Lester (nlester@umail.ucsb.edu)
Fermín Moscoso del Prado Martín (fmoscoso@linguistics.ucsb.edu)
Department of Linguistics
University of California, Santa Barbara, South Hall 3432
Santa Barbara, CA 93106 USA

Abstract

Does syntactic information affect the production of bare nouns? Research into this issue has explored word-specific features (e.g., gender). However, word-independent syntactic distributions may also play a role. For example, studies of word recognition have uncovered strong effects of the diversity of a word's syntactic distribution – its syntactic flexibility – on response times in the lexical decision paradigm. By contrast, studies of sentence production have produced strong but conflicted effects of syntactic flexibility. We propose that syntactic flexibility also affects production of individual words. We reanalyze a database of previously collected timed picture naming data using two novel measures of syntactic flexibility, one based on the relations stemming from the noun, and one based on the relations extending to the noun. Our results show that nouns that project a diverse array of structures are produced faster, and those that are integrated into a diverse array of structures are produced slower.

Keywords: syntactic flexibility; word production; picture naming; entropy

Introduction

Language production is fast – very fast – but the speed of the system is not without variability. One source of variability is the diversity of options available for encoding a given message. Most mainstream models of speech production rely on the notion of interactive competition, whereby the presence of alternatives delays encoding (so-called *interference effects;* e.g., Levelt, Roelofs, & Meyer, 1999). Such effects have been observed at the level of word production with respect to shared syntactic features of competitors: It takes longer to produce a word if other syntactic competitors are simultaneously active. Some studies have investigated similar issues at the sentence level. In one study, the number of sentence structures allowed by a given arrangement of words – their *syntactic flexibility* – was shown to correlate negatively with speech onset latencies (V. S. Ferreira, 1996). However, more recent research has reported the opposite findings: greater syntactic flexibility produced traditional interference effects, with slower production onsets for clauses with greater syntactic flexibility (Myachykov et al., 2013).

The clause-level studies have focused on the arrangement of arguments into abstract syntactic frames. One question that has not been addressed is how the syntactic flexibility of nouns affects their accessibility in production. Much work has been devoted to outlining the representation of syntactically relevant features of nouns, such as gender, as well as the time course of their activation across different sorts of tasks. However, to our knowledge, no study has specifically addressed whether a noun's retrieval depends on the diversity of its syntactic distributions. Are nouns that exhibit greater syntactic flexibility easier or harder to produce?

Syntactic Aspects of Lexical Access

All models of word production posit associations between individual lexical items and syntactic information (Caramazza, 1997; Dell, 1986; Levelt, Roelofs, & Meyer, 1999). However, they differ with respect to the time-course and directionality of activation between word-level and syntax-level nodes. For the most part, this research has focused on whether words share access to abstract specifications of gender and whether these gender-specifying nodes feed back to the lexical nodes during lexical retrieval. These studies typically rely on the picture-word interference task. In this task, participants are asked to name an image while ignoring a distractor word. The word either matches or conflicts with the target name with respect to its gender. The results have been split. If participants are asked to produce the picture name in a syntactic frame that requires access to the noun's gender (e.g., *die Katze*), then targets with matched-gender distractors are produced slower than mis-matched pairs (e.g., Schriefers & Teruel, 2000). Effects of this kind are typically attributed to competition between the distractor and target for access to a shared gender node. By contrast, when participants are asked to name a picture with a bare noun in the same task (e.g., *Katze*), then no interference effect is observed in the matched condition (La Heij, Mark, Sander, & Willeboorsde, 1998). This has led some to conclude that syntactic features are accessed preferentially during syntactic (i.e. phrasal), as opposed to purely lexical, production. However, a small number of studies eliciting bare-noun names in Romance languages have uncovered interference effects in the gender-matched condition (e.g., Cubelli et al., 2005). Thus, there is some support for activation of grammatical features in the picture-naming task, regardless of whether the form of the response requires that information.

Further support for universal access of syntactic features in bare-noun picture naming comes from a similar phenomenon – the mass/count distinction in English nouns (Gregory, Varley, & Herbert, 2012). In English, the mass/count noun distinction controls determiner type. For example, in their singular form, count nouns such as *cat* can take the determiner *each* (*each cat*), while mass nouns such

as *flour* cannot (?**each flour*). Gregory et al. (2012) found that bare-noun picture naming of count and mass nouns was facilitated by a grammatically congruent determiner prime. They attribute this finding to pre-activation of the syntactic feature shared between the determiner prime and the target word.

So far, we have seen that grammatical classes are relevant for word processing in production. Less clear is the extent to which syntactic relations – the scaffolds into which words are fit – impact lexical retrieval. To make this contrast more clear, consider that grammatical classes such as gender have consequences for the formal properties of words and their dependents (e.g., different patterns of agreement). However, they do not constrain the syntactic distributions of words. For example, the German words *Katze* and *Hund* belong to two different gender classes which take different forms of the definite determiner (*die* vs. *der*), along with different patterns of agreement elsewhere in the grammar. Nevertheless, both can occur with definite determiners, whatever their form. Going further, both may occur with adjectives, as subjects of sentences, as objects of prepositions, and so on. Crucially, gender does not dictate the availability of these relations, only the forms of the relata[1]. If information about grammatical class is activated when we access nouns for production, even when it doesn't impact the form of the utterance, then it is possible that other abstract syntactic information is likewise activated. Some recent research from comprehension supports this possibility.

Baayen et al. (2011) showed that the shape of a noun's (partially lexicalized) syntactic distribution impacts the time it takes to recognize that noun in visual lexical decision. They collected all trigrams of the form *preposition + determiner + noun* (for example, *on the table*). They then measured the diversity of the prepositions occurring with each noun and correlated this measure with lexical decision latencies. Nouns that occurred more evenly across a wider array of prepositions were recognized faster than words that skewed toward a limited array of prepositions. While interesting, these results speak only indirectly to the question raised above. This is because the trigrams used in the study do not contrast syntactic types, but rather lexical types (prepositions) within a single syntactic relation (prepositional object). Moreover, the proximity of the preposition and the noun opens the possibility that the facilitatory effect might actually be semantic in nature. Variability in small-scale co-occurrence windows is often assumed to reflect or indeed to constitute the meaning of words (Bullinaria & Levy, 2007).

In a follow-up study, Lester and Moscoso del Prado Martín (2015) compared the distributions of nouns across prepositional phrases within a parsed corpus of English writing. Crucially, the distance between noun and preposition could be of any size. First, they replicated the negative correlation between diversity and response time. To investigate the source of the effect in more detail, they compared the syntactically defined diversity space with a semantically-defined space (based on Latent Semantic Analysis vectors). They found almost zero correlation between the two. Second, they found that words with similar distributions in the syntactic space did not prime one another in an overt semantic priming task, while words with similar distributions in the semantic space did. Third, they showed that increasing average distance between the nouns and their prepositions reduced the strength of the facilitation, but did not eliminate or overturn it. That is, small-scale co-occurrence does seem to contribute something to our measure: when nouns tend to occur closer the prepositions, the facilitation is increased. However, it alone cannot explain the continued facilitation at longer distances. These findings suggest that something other than meaning must be contributing to the facilitation. A natural candidate is syntax.

These studies show that words encode abstract distributional information about their use and that this information becomes active in comprehension, even for isolated words. However, no study has yet compared syntactic types; the two studies above only treat the relation between a preposition and its object. Consequently, no study has measured diversity at a purely syntactic level, that is, without regard for the particular lexical types that anchor that relation opposite the target noun. Recall that the operationalizations of diversity outlined above included explicit reference to particular prepositions (e.g., how many times *table* occurs with *on* vs. *near* vs. *at* vs. ...). Finally, no study has yet examined how such effects play out in word production.

In the present study, we address these shortcomings by considering how flexibility across the entire set of grammatical relations afforded by a noun affects onset latencies in picture naming. We refer to this measure as the *total flexibility*. We further use the functional contrast between syntactic heads and dependents to decompose the total syntactic flexibility. On the one hand, we take the diversity of syntactic relations through which a noun serves as a dependent to some other head. Essentially, this measure captures how easily a noun can be integrated into a syntactic structure. We refer to this measure as *dependent flexibility*. On the other hand, we take the diversity of syntactic relations for which a noun can serve as head. This measure captures the structure-building potential of the noun. We refer to this measure as *head flexibility.* Importantly, all three of our measures are taken over abstract syntactic relations; they make no reference to any words besides the targets.

[1]Other grammatical classes may constrain the syntactic possibilities of nouns. For example, English mass nouns like *water* may function as their own phrase in non-metalinguistic contexts (*Water is wet*) while count nouns typically cannot (**Cat is furry*). However, these constraints are likely to be relatively minor. As the examples above demonstrate, syntactic distributions are at least partially independent of grammatical class.

Hypotheses

We contrast two general predictions regarding the behavior of our measures. We base these predictions on findings from sentence production as no comparable data exist for lexical access.

— The opportunistic hypothesis: Nouns with greater syntactic flexibility will be produced faster.

The first possibility derives from what Myachykov et al. (2013) termed the *opportunistic hypothesis*. The opportunistic hypothesis states that syntactic flexibility is generally facilitative. The assumption is that production operates most smoothly when the speaker has the most options for completing the current syntactic scaffold. Therefore, syntactically diverse nouns should be produced faster than syntactically limited nouns. Indirect support for this hypothesis comes from research on sentence production, which has shown that people will start speaking sooner when they have more syntactic options available for translating a scrambled word list into a grammatical sentence (V. S. Ferreira, 1996).

— The strategist hypothesis: Nouns with greater syntactic flexibility will be produced slower.

The second possibility derives from what we refer to as the *strategist hypothesis*. The strategist hypothesis states that more options require more deliberation before a choice can be made. Therefore, syntactically diverse nouns should be produced slower than syntactically limited nouns. Support for this possibility again comes from the sentence production literature. Myachykov et al. (2013) had speakers of Russian and English speakers describe scenes depicting transitive events. They found that Russian speakers took longer to initiate their descriptions than the English speakers. The authors attribute the finding to the much larger number of possible syntactic constructions available to Russian speakers (12 in total, as opposed to 2 for English). Hwang and Kaiser (2014), using a task similar to Ferreira (1996), likewise found slower onset latencies for syntactically flexible stimuli in Korean.

Neither hypothesis treated so far contrasts the effects of head and dependent flexibility. This is because no study to our knowledge has investigated whether different syntactic functions may respond differently to changes in diversity. This, coupled with the conflicted nature of the findings from sentence production, leads us to eschew any specific predictions with respect to our head and dependent measures.

Methods

To evaluate these hypotheses, we reanalyze the mean speech onset latencies from a large-scale picture-naming study (Bates et al., 2003; since published as part of the *International Picture Naming Project*). These data are ideal for testing the role of syntactic distributions in lexical access because they were collected using the bare-noun naming paradigm. This means that any effect of syntax that we observe cannot be attribute to features of the task (e.g., the fact that a speaker was required to activate a particular syntactic scaffold when responding).

Data

Bates and colleagues showed speakers of seven different languages a series of black-and-white line drawings and asked them to provide a bare-noun label (e.g., *cat*) for each image. They were instructed to say the name as quickly and fluently as possible. This instruction is critical: Ferreira and Swets (2002) show that speeded production tasks can induce more opportunistic processing strategies. Therefore, the data may be biased towards confirming the opportunistic hypothesis. Reaction times and errors were collected.

From the full data set, we extracted the mean onset RTs for the dominant, or majority names offered for each image. Of the 520 dominant responses, we took the subset of non-phrasal forms by omitting names that standard English orthography would divide into multiple words (such as *fishing pole*). We estimate the syntactic diversity of the remaining words on the basis of their distributions within the filtered Charniak-parsed treebank of the Open American National Corpus (OANC). This treebank contains phrase-structure parses of approximately 363,000 sentences of written American English from several genres, totaling around 11M words. We retrieved all parse trees containing the target words so long as the latter were tagged as common singular nouns We observed 416 of the original 520 word types in this sample (frequency in OANC: mean = 137.2; median = 42).

To tease apart the effects of flexibility for nouns as heads and nouns as dependents, we converted the phrase-structure trees that we culled from the OANC into *dependency graphs* using the Stanford dependency parser (de Marneffe, MacCartney, & Manning, 2006). In the dependency graph formalism, grammatical relationships are represented as ternary units consisting of a *head*, a *dependent*, and a typed *relation* (Tesnière, 1959). For example, *the fat cat* in the sentence *The fat cat ate the tuna* contains two dependencies: (1) a determiner relation *det* headed by *cat* and projected onto *the* and (2) an adjectival-modification relation *amod* headed by cat and projected onto *fat.* In the Stanford notation, these relationships surface as det(cat-3, the-1) and amod(cat-3, fat-2), respectively, where the numbers indicate sequential positioning in the clause. This formalism allows us to implement our measures in a simple, straightforward way. Dependent flexibility is operationalized as the diversity of relations for which the target noun is a dependent. Examples include the *nsubj* relation, which links a noun to a verb as its subject, or the *pobj* relation, which links a noun to a preposition as its object. Head flexibility is operationalized as the diversity of relations for which the target noun serves as head. Examples include the *amod* relation, which links an adjective to a noun, or *det*, which links a determiner to a noun. Importantly, we only care

about the target nouns and the relations; we ignore the words that are linked to the targets via the relation.

Syntactic Measures

To measure the diversity of the relation types, we generated three sparse matrices. Each matrix captures the frequency distributions of the target nouns (rows) across the set of relation types (columns). We constructed one such matrix for dependent relations, one for head relations, and one for all relations, regardless of whether the target was head or dependent.

For each word vector in each matrix, we computed the *Shannon entropy* (Shannon, 1948) of the distribution. This measure is particularly good at capturing diversity as it considers not only the number of types of relations, but also the relative frequencies of those types. The measure thus increases both with the number of types and with a more even spread of the instances across those types. Shannon entropy, denoted *H*, is defined as the mean of the weighted negative log probabilities of event *x* taken across outcomes *i* (Eq. 1).

$$H = -\sum_{i=1} p(x_i) \log\ p(x_i) \qquad (1)$$

Eq. 1 reflects the *maximum-likelihood* estimate of diversity within our vectors. As is, this measure will necessarily underestimate the true diversity of the sample. This is because we cannot be sure that the observed set of relations actually exhausts the set of possible relations for our targets. To counter this downward bias, we smooth our entropy estimates using the Chao, Wang, & Jost (2013) method (see Moscoso del Prado Martín, in press for a demonstration of the reliability of this message for word frequency distributions). We abbreviate our smoothed measure of head flexibility H_h and our smoothed measure of dependent flexibility as H_d. We abbreviate our baseline total flexibility measure as H_t.

Additional Controls

We also gathered information about the sentential distributions and semantic properties of each noun attested in our sample to avoid potential non-syntactic confounds. A word's typical positioning within a sentence may reflect aspects of its accessibility. Most models of sentence production hold that word order is partially determined by the relative speed with which the words are accessed (within the syntactic constraints afforded by the grammar of the language; Gletiman et al., 2007). Therefore, a word's general likelihood of showing up earlier in clauses might translate into faster bare-noun production in picture naming. In particular, words that frequently occur prior to the verb (e.g., as sentential subject or left-dislocated topic) might simply be more discourse-salient for speakers (better *starting points*; Konopka & Brown-Schmidt, 2014). Thus, we calculate the average relative positioning (position over sentence length) of each word, as well as the likelihood of the word preceding the sentential *root* dependency (instantiated lexically by the main verb).

Semantic effects in the timed picture naming tasks are split. Bormann (2011) found no difference in onset RTs for words from large or small semantic neighborhoods. However, Bates et al. (2003) report strong interference effects for onset RTs when the picture being described elicited more different names across participants. We likewise include the number of other names provided as a control. We also include a predictor of referent animacy, as animates are well known to be more accessible than inanimates (e.g., Ariel, 1991).

Finally, we include a number of other controls known to affect picture naming latencies: (log) lemma word frequency (CELEX), phonological complexity, visual complexity of the stimulus image, and subjective ratings of conceptual complexity.

Results

We computed several generalized additive models (GAMs) predicting mean (non-transformed) picture-naming RTs as dependent variable. As a first step, we added all control predictors (with spline-based smooths for all numeric predictors). After these, we added our total-flexibility measure H_t (with and without smoothing). As expected, we found a facilitatory linear effect of word frequency (β = -33.81, p<.0001) and a facilitatory effect of Bates and colleagues' measure of name agreement (p<.0001). Our baseline flexibility measure H_t had no effect (with or without smoothing; p>0.2).

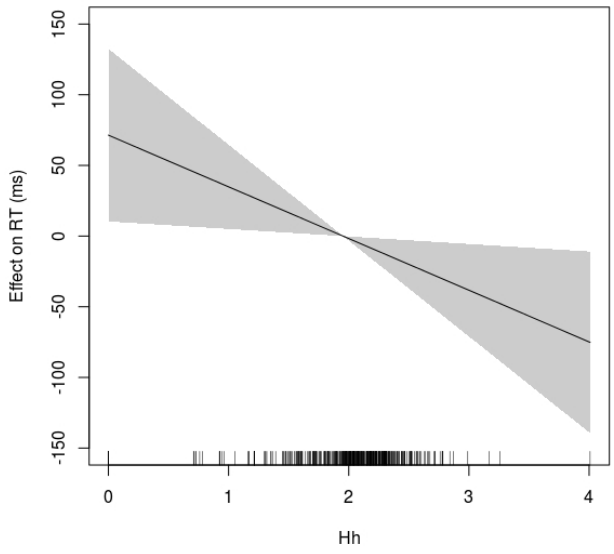

Figure 1: Facilitatory effect of H_h

In the next round, we repeated the model selection process, but substituted our measures of head flexibility H_h and dependent flexibility H_d for the total measure H_t. Collinearity between the head and dependent measures was

within the acceptable range (c=10.37)[2] as was the collinearity between each and word frequency (c<10 for both). The model returned significant effects of the same two control variables at essentially identical strengths. It also uncovered a facilitatory effect of H_h (p<.05; Figure 1) and an inhibitory effect of H_d (p<.05; Figure 2). Both effects are of roughly equal size, with a difference in extremes of ~100 ms.

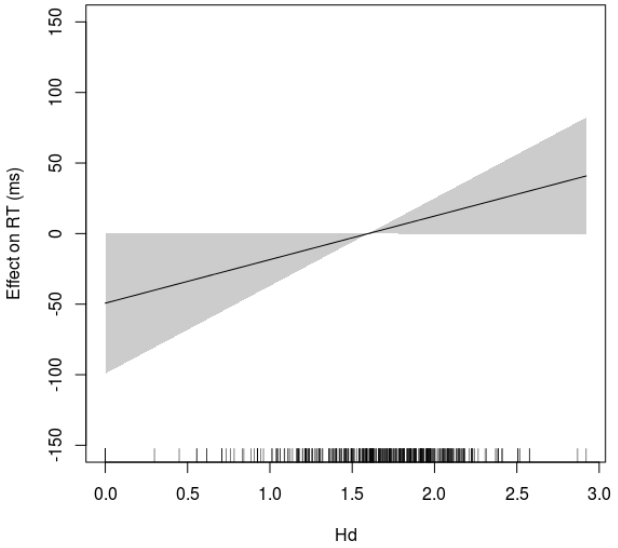

Figure 2: Inhibitory effect of H_d

Discussion

Our results provide first evidence that a noun's syntactic distribution affects its accessibility in word production. These findings support and expand prior work which has demonstrated that speakers access syntactic features when producing words in isolation (Cubelli et al., 2005; Gregory et al., 2012). Moreover, these findings suggest that words encode quite fine-grained information about their syntactic potential. Finally, we show that the function of the noun within syntactic relations plays a critical role in determining its accessibility. We find facilitation for nouns that are heads of a diverse array of relations and inhibition for those that are dependents of a diverse array of relations.

The facilitatory effect of head diversity is compatible with the opportunistic hypothesis, which states that production proceeds best when we have more options for structural encoding (V. S. Ferreira, 1996). Following this logic, nouns that provide more possibilities for expanding their phrasal nodes could be accessed more quickly because their selection does not immediately commit the speaker to any particular phrasal structure. Of course, the utterances analyzed here consist of single words. However, the opportunistic hypothesis makes reference only to *possible* structures. In that sense, the speed of lexical access need not depend on the intention to build a broader syntactic unit around the word. Rather, it may depend on the array of structures suggested by its activation. But this is only part of the story. Recall that several studies have shown that other forms of syntactic information, such as gender, can affect production of single words (e.g., Cubelli et al., 2005). Taken together, these facts suggest that the bare-noun naming paradigm may not be effective at eliminating syntactic processing, at least not entirely. This makes sense if we consider that people rarely produce words apart from a syntactic context. If when the system activates a word, it expects strongly to produce that word in some syntactic arrangement, then it may persist in activating syntax-relevant information despite the constraints imposed by the task.

Turning to dependent flexibility, we uncovered an inhibition effect: nouns that showed a greater diversity of embedding contexts were produced slower. This finding is compatible with the strategist hypothesis (Myachykov et al., 2013), which predicts that more options make it more difficult to settle on the final outcome. This result may be directly related to the task. The participants were required to produce only a single word. The single-word utterance is one of many syntactic templates in which a word could have been experienced (e.g., in exchanges like A. *What did you say?* B. *Cat.*). In dependency grammar, this would be represented as the ROOT relation (a special 'headless' dependency type that only takes a dependent). Thus, selecting the ROOT relation from among the other possible structures may have resulted in a traditional interference effect. This point distinguishes head and dependent diversity within the context of bare-noun picture naming. When producing a noun in the ROOT relation, one has actually committed to a syntactic structure – one of many available to call the noun as dependent. However, when there is no intention to elaborate the noun into a noun phrase, the diversity of its possible head relations can facilitate access without interference. The structures that would otherwise be competing for selection now support access by feeding activation back into the target lexeme. The more diverse the feedback network, the faster the access.

This pattern of results underscores two critical points regarding language processing: (1) words are finely articulated syntactic entities whose history of use partially determines how efficiently they are processed – even in isolation! – and (2) flexibility is not a simple component of syntactic entities, but one that interacts with different functional domains to help or hinder processing. Perhaps words and syntactic structures are much more tightly linked than is typically acknowledged. If so, then the effects of flexibility observed here might be attributed to the flow of

[2] Collinearity between our three flexibility measures exceeded the acceptable range (c=24.5). To clean this collinearity, we performed a principal components analysis (PCA) on the three measures. The PCA yielded a component that distinguished H_h (.64) from H_m (-.76) with near 0 loading for H_t. When we entered this component into the GAM frame described above, we found a noisier but significant effect ultimately equivalent to the individual effects of H_h and H_m (p < .05). However, to maximize the clarity of results attributable to our decomposed measures, we report the analysis that simply omits H_t.

activation within the network – top-down from the syntactic tier or bottom-up from individual words. This type of processing model would be more in line with current linguistic models of representation, for which grammar and lexicon form a continuum (e.g., Goldberg, 2006). Architecturally, it wouldn't be a matter of representing syntax *in* the noun so much as *with* the noun. At any rate, continued exploration and refinement of these diversity effects will no doubt sharpen our perspective on the relationship between words and structure.

References

Ariel , M. (1991). The function of accessibility in a theory of grammar. *Journal of Pragmatics*, *16*, 443-463.

Baayen, R. H., Milin, P., Filipović-Đurđević, D., Hendrix, P. & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. Psychological Review, 118, 438-482.

Bates, E., D'Amico, S., Jacobsen, T., Székely, A., Andonova, E., Devescovi, A., Herron, D., Lu, C. C., Pechmann, T., Pléh, C., Wicha, N., Federmeier, L. Gerdjikova, I., Gutierrez, G., Hung, D., Hsu, J., Iyer, G., Kohnert, K., Mehotcheva, T., Orozco-Figueroa, A., Tzeng, A., & Tzeng, O. (2003). Timed picture naming in seven languages. *Pscyhonomic Bulletin and Review*, *10*, 344-380.

Bock, J. K. & Ferreira, V. S. (2014). Syntactically speaking. In M. Goldrick, V. S. Ferreira, & M. Miozzo (Eds), The Oxford Handbook of Language Production (pp. 21-46). New York: Oxford University Press.

Bormann, T. (2011). The role of lexical-semantic neighborhood in object naming: Implications for models of lexical access. *Frontiers in Psychology*, *2*. 1-11.

Bullinaria, J. A. & Levy, J. P. (2007). Extracting Semantic Representations from Word Co-occurrence Statistics: A Computational Study. *Behavior Research Methods*, *39*, 510-526.

Caramazza, A. (1997). How many levels of processing are there in lexical access? *Cognitive Neuropsychology*, *14*, 177–208.

Chao, A., Wang, Y. T., & Jost, L. (2013). Entropy and the species accumulation curve: a novel estimator of entropy via discovery rates of new species. *Methods in Ecology and Evolution*, *4*, 1091–1110

Cubelli, R., Lotto, L., Paolieri, D., Girelli, M., & Job, R. (2005). Grammatical gender is selected in bare noun production: Evidence from the picture-word interference paradigm. *Journal of Memory and Language*, *53*, 42–59.

Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review*, *93*, 283-321.

de Marneffe, M-C., MacCartney, B., & Manning, C. D. (2006). Generating Typed Dependency Parses from Phrase Structure Parses. In *LREC 2006*.

Ferreira, F., & Swets, B. (2002). How incremental is language production? Evidence from the production of utterances requiring the computation of arithmetic sums. *Journal of Memory and Language*, *46*, 57–84.

Ferreira, V. S. (1996). Is it better to give than to donate? Syntactic flexibility in language production. *Journal of Memory and Language*, *35*, 724–755.

Gleitman, L., January, D., Nappa, R., & Trueswell, J.C. (2007). On the give and take between event apprehension and utterance formulation. *Journal of Memory and Language*, *57*, 544-569.

Goldberg, A. E. (2006). *Constructions at work*: *The nature of generalization in language.* Oxford: Oxford University Press.

Gregory, E., Varley, R., Herbert, R. (2012). Determiner primes as facilitators of lexical retrieval. *Journal of Psycholinguistic Research*, *41*, 439-453.

Hwang, H., & Kaiser, E. (2014). Having a syntactic choice is not always better: The effects of syntactic flexibility on Korean production. *Language, Cognition, and Neuroscience*, *29*, 1115-1131.

Konopka, A. E., & Brown-Schmidt, S. (2014). Message encoding. In M. Goldrick, V. S. Ferreira, & M. Miozzo (Eds), The Oxford Handbook of Language Production (pp. 3-20). New York: Oxford University Press.

La Heij, W., Mark, P., Sander, J., & Willeboorsde, E. (1998). The gender congruency effect in picture word task. *Psychological Research*, *61*, 209–219.

Langacker, R. W. (1987). *Foundations of Cognitive Grammar: Theoretical prerequisites.* Stanford: CSLI.

Lester, N. A., & Moscoso del Prado Martín, F. (2015). Constructional paradigms affect visual lexical decision latencies in English. In Noelle, D. C., Dale, R., Warlaumont, A. S., Yoshimi, J., Matlock, T., Jennings, C. D., & Maglio, P. P. (Eds.), *Proceedings of the 37th Annual Meeting of the Cognitive Science Society* (pp. 1320-1325). Austin, TX: Cognitive Science Society.

Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, *22*, 1–75.

Moscoso del Prado Martín, F. (in press). Vocabulary, grammar, sex, and aging. *Cognitive Science.*

Myachykov, A., Scheepers, C., Garrod, S., Thompson, D., & Fedorova, O. (2013). Syntactic flexibility and competition in sentence production: The case of English and Russian. *The Quarterly Journal of Experimental Psychology*, *66*, 1601-1619.

Schriefers, H., & Teruel, E. (2000). Grammatical gender in noun phrase production: The gender interference effect in German. *Journal of Experimental Psychology: Learning memory, and Cognition*, *26*, 1368-1377.

Shannon, C. E. (1948). A mathematical theory of communication. The Bell System Technical Journal, 28(3), 379-423.

Tesnière, L. (1959). Éléments de syntaxe structurale. Paris: Klincksieck.