

Incremental semantic restriction and subjectivity–based adjective ordering¹

Gregory SCONTRAS — *University of California, Irvine*

Galia BAR-SEVER — *Pomona College*

Zeinab KACHAKECHE — *University of California, Irvine*

Cesar MANUEL ROSALES JR. — *University of California, Irvine*

Suttera SAMONTE — *San Francisco State University*

Abstract. Adjective ordering preferences (e.g., *big blue box* vs. *blue big box*) are robustly attested in many unrelated languages (Dixon, 1982). Scontras et al. (2017) showed that adjective subjectivity is a robust predictor of ordering preferences in English: less subjective adjectives occur closer to the modified noun. In a follow-up to this finding, several authors have claimed that pressures from successful reference resolution and the hierarchical structure of modification explain subjectivity-based ordering preferences (Simonič, 2018; Franke et al., 2019; Scontras et al., 2019). In cases of restrictive modification, adjectives that compose with the nominal later will classify a smaller set of potential referents (e.g., the set of boxes vs. the set of blue boxes). To avoid alignment errors where a listener might mis-characterize the intended referent, speakers introduce the more error-prone (i.e., more subjective) adjectives later in the hierarchical construction of nominal structure; the structure linearizes such that subjectivity decreases the closer you get to the modified noun. The current study explores the predictions of this reference-resolution story by examining adjective ordering cross-linguistically: when adjectives incrementally restrict a nominal denotation, there should be pressure toward subjectivity-based orderings, but, in the absence of incremental restriction, such pressures should not obtain.

Keywords: adjective ordering, subjectivity, hierarchical structure, conjunction, Arabic, English, Spanish, Tagalog.

1. Introduction

Adjective ordering preferences influence the relative order of adjectives in multi-adjective strings, for example *big blue box* vs. *blue big box*. English speakers exhibit a robust preference for the former order, such that color adjectives are preferred closer to the modified noun than size adjectives, and this same preference has been reported in a host of unrelated languages (e.g., Dixon, 1982; Sproat and Shih, 1991). Various proposals have been advanced to account for these preferences, from articulated syntactic hierarchies (Cinque, 1994; Scott, 2002) to appeals to psychological accessibility (Whorf, 1945; Martin, 1969) or ease of parsing (Bever, 1970). Recently, proposals that focus on issues of adjective meaning have gained large-scale empirical support.

Scontras et al. (2017) used corpus and behavioral data to show that adjective subjectivity is a robust predictor of adjective ordering preferences in English, such that less subjective adjectives are preferred closer to the modified noun. In *big blue box*, speakers perceive *blue* as less subjective than *big*, and so *blue* occurs closer to the noun.

In an attempt to explain this robust empirical generalization—that subjectivity predicts adjective ordering preferences—several authors have arrived at the conclusion that ordering ad-

¹We thank the audience and organizers of Sinn und Bedeutung 24 in Osnabrück. Thanks also to members of the Quantitative Language Science Collective at UC Irvine for feedback on earlier versions of this work.

jectives with respect to decreasing subjectivity maximizes communicative success (Simonič, 2018; Hahn et al., 2018; Franke et al., 2019; Scontras et al., 2019). While different authors use different assumptions in their accounts, most agree that pressures from successful reference resolution and the hierarchical structure of nominal modification stand to explain subjectivity-based ordering preferences (Simonič, 2018; Franke et al., 2019; Scontras et al., 2019). In cases of restrictive modification, adjectives that compose with the nominal later will classify a smaller set of potential referents (e.g., the set of boxes vs. the set of blue boxes). To avoid alignment errors where a listener might mis-characterize the intended referent, speakers introduce the more error-prone (i.e., more subjective) adjectives later in the hierarchical construction of nominal structure; the structure linearizes such that subjectivity decreases the closer you get to the modified noun.

The current study explores the predictions of this communicative-efficiency story by examining adjective ordering cross-linguistically: when adjectives incrementally restrict a nominal denotation, there should be pressure toward subjectivity-based orderings, but, in the absence of incremental restriction, such pressures should not obtain. We investigate adjective ordering in both pre-nominal languages where adjectives precede nouns (English, Tagalog) and post-nominal languages where adjectives follow nouns (Spanish, Arabic). We also explore the role of linking particles that mediate the composition of modifiers (Rubin, 1994; Scontras and Nicolae, 2014), contrasting such particles with run-of-the-mill conjunction. Our results further confirm the empirical generalization concerning subjectivity in adjective ordering; the results also add new support for accounts of subjectivity in adjective ordering that rely on incremental semantic restriction that tracks the hierarchical structure of modification.

2. Subjectivity-based ordering preferences

We begin by reviewing the empirical methodology of Scontras et al. (2017), which will serve as the foundation for our empirical investigations. We then review the details of Franke et al.'s (2019) proposal concerning the role of subjectivity in adjective ordering, identifying two predictions made by this proposal.

2.1. Evidence from English

Scontras et al. (2017) investigated whether aspects of adjective meaning explain adjective ordering preferences. Specifically, Scontras et al. tested whether adjectives are ordered with respect to decreasing subjectivity, such that adjectives perceived to be less subjective are preferred closer to the modified noun (Hetzron, 1978; Tucker, 1998; Hill, 2012). The authors began by measuring ordering preferences for 26 relatively frequent adjectives from seven semantic classes. Experimental participants were presented with a series of adjective-adjective-noun pairs that differed in the relative order of the adjectives, and they adjusted a slider to indicate which ordering sounded more natural (e.g., *metal tiny chair* vs. *tiny metal chair* in Figure 1). To validate this behavioral measure, Scontras et al. conducted a corpus analysis of naturally-occurring multi-adjective strings. The behavioral and corpus measures were found to be highly correlated ($r^2 = 0.83$), so the authors concluded that the behavioral measure effectively captured ordering preferences that speakers use when they form multi-adjective strings.

To measure subjectivity, Scontras et al. used a faultless disagreement task (Kölbel, 2004; Mac-

Progress:

Which description of the chair sounds more natural?

"the metal tiny chair" "the tiny metal chair"

Adjust the slider to indicate your preference.

Figure 1: Sample trial from the ordering preferences experiment of Scontras et al. (2017).

Farlane, 2014). Participants were presented with short dialogues in which two speakers disagreed about a property ascription (e.g., whether some cheese was *rotten* in Figure 2). Participants had to decide whether the two speakers could both be right while disagreeing (i.e., whether they could faultlessly disagree), or whether one of the speakers must be wrong. Scontras et al. used an adjective's potential for faultless disagreement as an index of adjective subjectivity.²

To evaluate the subjectivity hypothesis, Scontras et al. compared their subjectivity scores with the ordering preferences they measured. Subjectivity was found to explain between 85% and 88% of the variance in the ordering preferences. To test the generalizability of their findings, the authors also looked at ordering preferences for 74 adjectives found to naturally occur in multi-adjective strings in the Switchboard corpus of English telephone conversations; subjectivity accounted for 61% of the variance in the ordering preferences for these adjectives. Thus, Scontras et al. found strong evidence in support of their hypothesis: an adjective's meaning does predict its distance from the noun it modifies, such that less subjective adjectives occur closer to the modified noun.

2.2. Subjectivity-based ordering maximizes communicative success

With clear evidence for the empirical generalization that subjectivity predicts adjective ordering preferences, the task turns next to explaining why subjectivity should play its role in adjective ordering. A number of proposals have recently been put forth, and, while they rely on different sets of assumptions, all of these proposals agree that adjectives are ordered with respect to decreasing subjectivity in an effort to maximize the communicative success of multi-adjective nominals (Simonič, 2018; Hahn et al., 2018; Franke et al., 2019; Scontras et al., 2019). In other words, multi-adjective strings ordered with respect to decreasing subjectivity are more likely to allow a speaker to successfully communicate their intended message to a listener. Here, we

²In a separate experiment, Scontras et al. measured subjectivity by asking participants how "subjective" a given adjective was. These raw subjectivity scores were found to be highly correlated with estimates of faultless disagreement ($r^2 = 0.91$). We use raw the subjectivity scores from Scontras et al. in our analysis of English conjunction below.

Progress:

Consider the following situation:

Eric and Kevin see the same cheese.

Eric says: "**That cheese is not rotten.**"

Kevin responds: "**You're wrong. That cheese is rotten.**"

Can both Eric and Kevin be right?

No, somebody must be wrong. Yes, it's a matter of opinion.

Figure 2: Sample trial from the faultless disagreement experiment of Scontras et al. (2017).

review the specific proposal advanced by Franke et al. (2019), pointing out places where its assumptions overlap with those of other accounts.

The starting observation for all of the recent accounts of subjectivity-based adjective ordering is that less subjective content is, in a sense to be specified, more useful for effectively and efficiently communicating about the world. When a listener hears the adjective *blue*, the set of potential referents they imagine is likely to be more constrained or less variable than the set of referents they imagine when hearing *big*; we find evidence of this divergence reflected in the different faultless disagreement scores assigned to the two adjectives. When determining the order of a multi-adjective string, the descriptive generalization is that speakers prefer to place more useful, less subjective content closer to the modified noun. To see why, we have to consider in more detail the adjectives' semantics.

Franke et al. (2019) focus on uses of adjectives that aid in establishing nominal reference.³ Thus, their aim is to make precise the notion that less subjective adjectives are more useful for successful reference resolution, and show that subjectivity-based adjective ordering leads to greater success. The authors assume the empirically-motivated context-dependent semantics from Schmidt et al. (2009). Under this semantics, an adjective like *big* characterizes those objects that meet the contextual cutoff for size. The cutoff is calculated on the basis of relative height by range, such that any object that falls within the top $k\%$ of the range of sizes in the context C counts as *big* in C . The corresponding lexical entry for *big* appears in (1), where $\text{size}(x)$ finds the size of some object x , max is the size of the largest object in C , min is the size of the smallest object in C , and $\theta = k/100$.

$$(1) \quad \llbracket \text{big} \rrbracket^C = \lambda x \in C. \text{size}(x) \geq (\text{max} - \theta \cdot (\text{max} - \text{min}))$$

Suppose the maximum object size in C is 10 (on some arbitrary scale), the minimum is 2, and we set the relevant threshold k at 50%; with these settings, the size cutoff for *big* would be 6, so that any object with size 6 or greater would count as *big* in C .

Franke et al. further assume that sequential adjectival modification, as in multi-adjective nom-

³This focus is shared by Simonič (2018) and Scontras et al. (2019).

inals like *big blue box*, can trigger sequentially intersective updates to the context.⁴ This move has adjectives farther from the noun interpreted with respect to contexts that have already been restricted by any closer adjectives. Thus, a phrase like *big blue box* can be interpreted as characterizing the set of blue boxes that count as big for the blue boxes in C . In (2), we schematically represent this incremental semantic restriction for multi-adjective nominals with two adjectives.

$$(2) \quad [[A_2 A_1 N]] = [[A_2]] [[A_1]]^{[M] \cap C}$$

For Franke et al., noise—what we described above in terms of differential communicative utility—enters at the level of perception as agents create their subjective representations of the discourse context.⁵ Whatever the true context, each agent arrives at their own representation. Crucially, more subjective properties (e.g., size) are assumed to more commonly lead to deviations between the true context and an agent’s representation. Because each conversational agent independently arrives at their own subjective representation of the context, more subjective properties more commonly lead to deviations between the two agents’ representations. Franke et al. propose that these deviations and our awareness of them contribute to an adjective’s perceived subjectivity.

With these assumptions in place, Franke et al. simulated 1,000,000 contexts (i.e., sets of potential referents) and recorded the probability of a listener correctly retrieving the intended referent on the basis subjectivity-based orderings (e.g., *big blue box*) vs. non-canonical reverse orderings (e.g., *blue big box*). By performing these comparisons, the authors compared two hypothetical groups of speakers, asking which has higher communicative success: the group using subjectivity-based orderings, or the group using the reverse orderings. The results clearly demonstrated that, on average, subjectivity-based orderings are more likely to lead to communicative success. Thus, the authors provide an evolutionary rationale for why natural languages prefer one ordering over another: subjectivity-based orderings are more conducive to our communicative aims. Crucially, this explanation does not require active reasoning comparing adjective subjectivity as speakers form their multi-adjective strings. Rather, certain orderings (i.e., the subjectivity-based ones) will turn out to be more useful to speakers in the long run, and so those orderings are more likely to stick around in the language. Language users will more commonly encounter those more useful subjectivity-based orderings, and so production can simply mirror the statistics of the input to yield robust subjectivity-based ordering preferences.

2.3. Predictions

We have a robust empirical generalization—subjectivity predicts adjective ordering preferences—and a proposal meant to explain the generalization—subjectivity-based adjective ordering maximizes communicative success. We now explore two predictions of this proposal, which we will test empirically below.

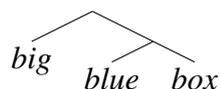
First, if communicative pressures deliver subjectivity-based ordering when adjectives incrementally restrict a nominal denotation, then we should find that these pressures apply whenever

⁴This assumption is shared with Simonič (2018) and Scontras et al. (2019).

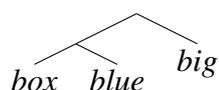
⁵Here is a point where the assumptions of Franke et al. deviate from those of other authors. Simonič (2018), Hahn et al. (2018), and Scontras et al. (2019) build the noise associated with subjectivity directly into the semantics of adjectives.

we have configurations of adjectives that result in incremental semantic restriction. Specifically, we should find subjectivity-based adjective ordering preferences in both pre-nominal languages like English, where adjectives precede the modified noun, and in post-nominal languages like Spanish or Arabic, where adjectives follow the noun. Such structures are given schematically in (3).

(3) a. *Pre-nominal structure:*



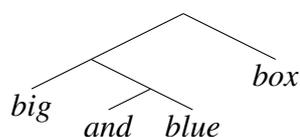
b. *Post-nominal structure:*



Recall the explanation from Franke et al. (2019) above: in cases of restrictive modification, adjectives that compose with the nominal later will classify a smaller set of potential referents (e.g., the set of boxes vs. the set of blue boxes); to avoid alignment errors where a listener might mis-characterize the intended referent, speakers introduce the more error-prone (i.e., more subjective) adjectives later in the hierarchical construction of nominal structure; the structure linearizes such that subjectivity decreases the closer you get to the modified noun. Crucially, the explanation applies regardless of whether the noun precedes or follows the noun. In (3), either configuration has adjectives farther from the noun composing later and therefore operating over a restricted set of potential referents. Thus, with either configuration, we predict that communicative pressures should interact with the hierarchical structure of multi-adjective modification to deliver subjectivity-based preferences.

Second, we predict that subjectivity-based preferences should not arise in the absence of incremental semantic restriction. Put differently, whenever the hierarchical structure is such that adjectives are not composing sequentially with the modified nominal, communicative pressure toward subjectivity-based ordering should not apply. One structure that disrupts incremental semantic restriction is multi-adjective strings formed via conjunction, as in (4).

(4) *Conjoined structure:*



With conjunction, the adjectives make their semantic contribution (i.e., restriction of the relevant context) together after they are conjoined, so pressures mediating the order in which adjectives make that contribution should not apply. Indeed, conjunction has been claimed to neutralize ordering preferences in English (e.g., Ford and Olson, 1975; Byrne, 1979), leading to the acceptability of otherwise unacceptable orderings. For example, while we might strongly disprefer *blue big box*, the same ordering formed via conjunction should be acceptable: *blue and big box*.

The remainder of the paper explores the predictions of this reference-resolution story for subjectivity-based preferences by examining adjective ordering cross-linguistically: when adjectives incrementally restrict a nominal denotation, there should be pressure toward subjectivity-based

orderings, but, in the absence of incremental restriction, such pressures should not obtain.

3. Testing the predictions

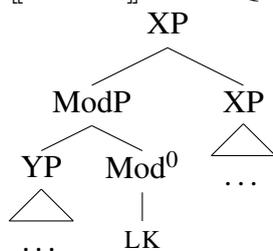
3.1. Tagalog

We begin with a look at Tagalog, a pre-nominal language where adjectives require a linking particle (*-ng/na*) to participate in modification structures (Foley, 1975; Rubin, 1994).⁶ An example multi-adjective nominal appears in (5); note the obligatory presence of LINKER.

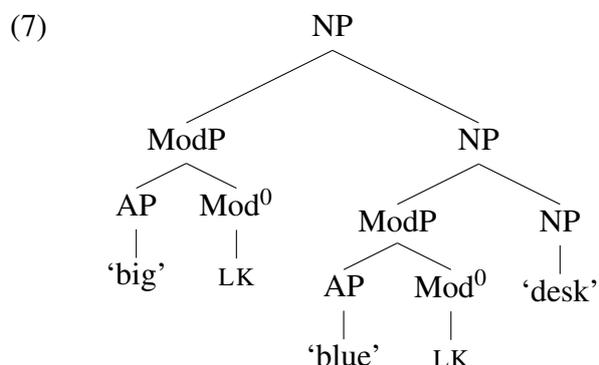
- (5) malaki-ng asul na mesa
 big-LK blue LK desk
 ‘big blue desk’

Tagalog’s LINKER appears in the presence of modification; some have analyzed the semantic contribution of LINKER similarly to that of conjunction, as in (6) (Rubin, 1994; Scontras and Nicolae, 2014).

- (6) a. $[[\text{LINKER}]] = \lambda P \lambda Q \lambda x. P(x) \wedge Q(x)$
 b.



As noted above, in English, conjunction has been claimed to neutralize adjective ordering preferences. We might therefore expect that Tagalog’s LINKER, with its conjunction-like semantics, also neutralizes ordering preferences. However, even if LINKER discharges the semantics of modification, adjectives still compose incrementally with the nominal. In (7), it is still the case that adjectives closer to the noun restrict the nominal denotation before adjectives that are farther away. If subjectivity-based preferences derive from incremental semantic restriction, these preferences should surface in Tagalog.



We therefore set out to determine (i) whether Tagalog possesses ordering preferences in the presence of LINKER, and, if so, (ii) whether subjectivity predicts those preferences.

⁶The material in this section summarizes the findings of Samonte and Scontras (2019).

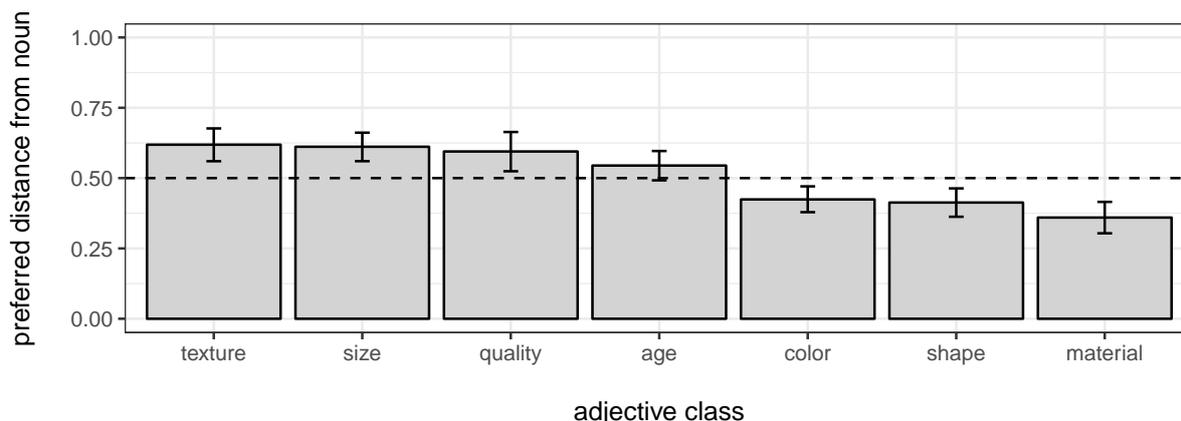


Figure 3: Tagalog naturalness ratings grouped by adjective semantic class. Higher values indicate that a class’s adjectives are preferred farther from the modified noun; lower values indicate that a class’s adjectives are preferred closer. The dashed line indicates chance level, or the absence of stable preferences. Error bars represent bootstrapped 95% confidence intervals drawn from 10,000 samples of the data.

3.1.1. Measuring preferences

We replicated Experiment 1: *Ordering preferences* from Scontras et al. (2017) using Tagalog translations of the original English materials.⁷

Participants. We recruited 90 participants through Amazon.com’s Mechanical Turk. On the basis of their responses to a post-test questionnaire, 24 Tagalog-speaking participants were identified; we include their data in the analyses reported below.

Procedure. Participants indicated their preferences for pairs of multi-adjective strings formed from 26 unique adjectives from seven semantic classes paired with ten nouns; the pairs differed on the relative order of the adjectives (e.g., *malaking asul na mesa* ‘big blue desk’ vs. *asul na malaking mesa* ‘blue big desk’). On each trial, adjectives and nouns were randomly chosen, with the constraint that the two adjectives were from different semantic classes. Participants completed a series of 26 trials.

Results. We averaged across participants’ ratings to calculate a single preferred-distance measure for each adjective; values ranged from 0 (always preferred closest to the noun) to 1 (always preferred farthest from the noun). Figure 3 plots the preferred-distance measures grouped by adjective class. All but the age adjectives deviate significantly from the random baseline (i.e., from 0.5), suggesting that Tagalog does indeed have stable preferences: some adjectives are reliably preferred closer to the noun, while others are reliably preferred farther away. Thus, we find strong evidence in support of stable ordering preferences in Tagalog, despite the obligatory LINKER.

⁷See Samonte and Scontras (2019) for the full details of the materials.

3.1.2. Measuring subjectivity

We next measured adjective subjectivity in Tagalog using a faultless disagreement task (cf. Expt. 1: *Faultless disagreement validation* from Scontras et al., 2017): to the extent that two speakers can be right while disagreeing about a property, the property admits that degree of faultless disagreement, which indexes adjective subjectivity.

Participants. We recruited 45 participants through Mechanical Turk; participants who had taken part in the ordering preferences experiment were not eligible. We identified eleven Tagalog speakers on the basis of their responses to a post-test questionnaire.

Procedure. Participants encountered a series of dialogues in which two speakers disagreed about a property description (e.g., whether or not some desk was ‘blue’). The task was to determine whether the two speakers could both be right while disagreeing (i.e., whether they could faultlessly disagree), or whether one speaker must be wrong. Participants completed a series of 26 trials, one for each of the adjectives tested in the ordering preferences experiment.

Results. Responses ranged from 0 (‘only one can be right’) to 1 (‘yes, it depends on what you believe’). For each adjective, we computed a mean faultless disagreement score by averaging across participants’ responses. We will use these subjectivity scores in the following subsection.

3.1.3. Comparing ordering preferences with subjectivity

With measures of the Tagalog ordering preferences and adjective subjectivity, we can ask whether subjectivity predicts those ordering preferences. Figure 4 plots ordering preferences against subjectivity scores for each of the 26 adjectives tested. There, we see that in Tagalog, as in English, subjectivity is a reliable predictor of individual adjective ordering preferences. Despite using diverging strategies to form modification structures, in both languages adjectives compose incrementally with the modified noun, which leads to subjectivity-based ordering preferences.

3.2. Spanish

We turn next to Spanish, where multi-adjective strings are post-nominal and commonly formed via conjunction (e.g., *el escritorio grande y azul* ‘the big and blue desk’).⁸ Some adjectives may occur pre-nominally; however, this strategy is not fully productive in the language. We therefore focused only on post-nominal multi-adjective strings. To investigate ordering preferences in Spanish, we replicated the methodology from Scontras et al. (2017) using Spanish translations of the English materials.⁹ Given that many speakers express a preference for multi-adjective strings formed via conjunction, we ran two separate ordering preferences experiments, with and without conjunction.

⁸Portions of the material in this section summarize the findings of Rosales Jr. and Scontras (2019); the conjunction-free results are novel.

⁹See Rosales Jr. and Scontras (2019) for full details.

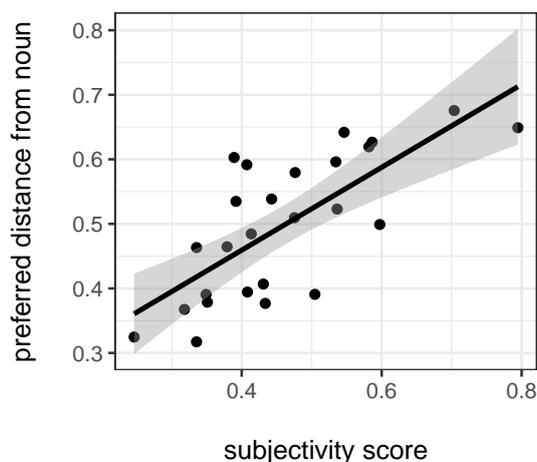


Figure 4: Tagalog ordering preferences plotted against subjectivity scores for each of the 26 adjectives tested. Subjectivity accounts for 54% of the variance in the ordering preferences ($r^2 = 0.54$, 95% CI = [0.22, 0.74]).

3.2.1. Measuring preferences with conjunction

We replicated the ordering preferences experiment described above for Tagalog, here using Spanish materials that always featured conjunction in the formation of multi-adjective strings.

Participants. We recruited 224 participants through Mechanical Turk; 48 were identified as native speakers of Spanish on the basis of a post-test questionnaire.

Procedure. The procedure was identical to the Tagalog ordering preferences experiment with the exception that adjectives were conjoined to form multi-adjective strings.

Results. Figure 5 plots average preferred-distance measures grouped by adjective class. We see that for all but one of the classes (i.e., the quality adjectives), participants did not provide systematic ratings that would evidence stable ordering preferences. In other words, from their responses we see that participants do not have clear preferences to place certain classes of adjectives closer or farther from the modified noun; we find a similar pattern at the level of individual adjectives.

3.2.2. Measuring preferences without conjunction

Although our consultants indicated a preference to use conjunction in multi-adjective strings, after publishing the initial results in Rosales Jr. and Scontras (2019), we learned that some speakers do not consider conjunction *necessary* for forming multi-adjective strings. Therefore, we repeated the Spanish ordering preferences experiment, this time without conjunction between the relevant adjectives.

Participants. We recruited 180 participants through Mechanical Turk, of which 22 were identified as native speakers of Spanish on the basis of a post-test questionnaire.

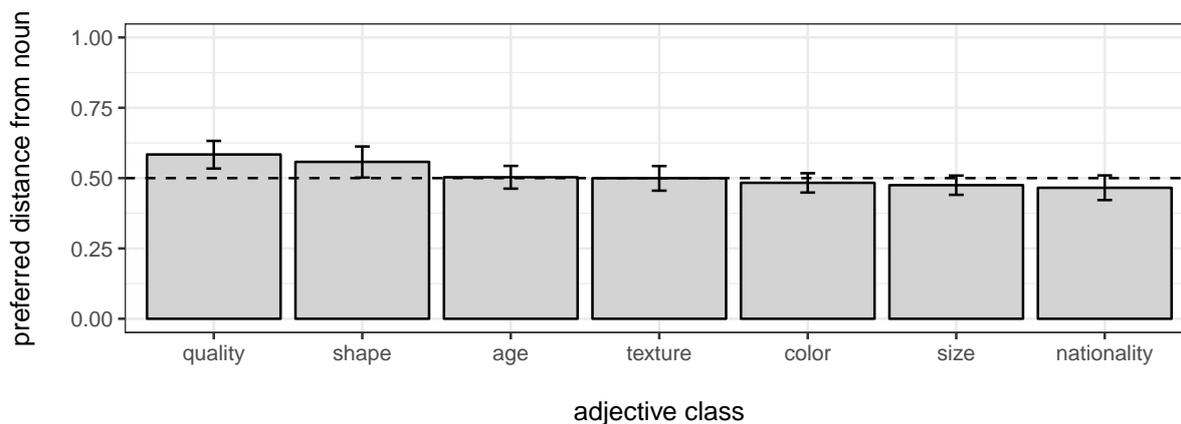


Figure 5: Spanish naturalness ratings for conjoined strings grouped by adjective semantic class.

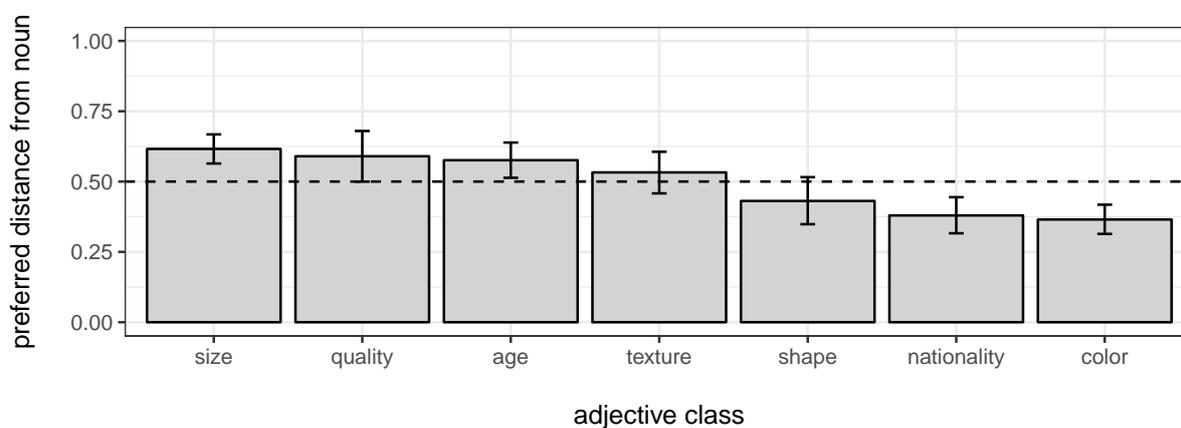


Figure 6: Spanish naturalness ratings for strings without conjunction grouped by adjective semantic class.

Procedure. The procedure was identical to the Spanish ordering preferences experiment above, with the exception that here we measured preferences for multi-adjective strings without conjunction.

Results. Figure 6 plots average preferred-distance measures. In contrast to the results in Figure 5, without conjunction we see clear evidence of ordering preferences in Spanish: certain classes are preferred closer to the noun (i.e., size, age), and others are preferred farther away (i.e., nationality, color).

3.2.3. Measuring subjectivity

To measure subjectivity, we replicated the faultless disagreement task using Spanish materials.

Participants. We recruited 106 participants through Mechanical Turk; 21 participants were identified as native speakers of Spanish.

Procedure. The procedure was identical to the faultless disagreement experiment described

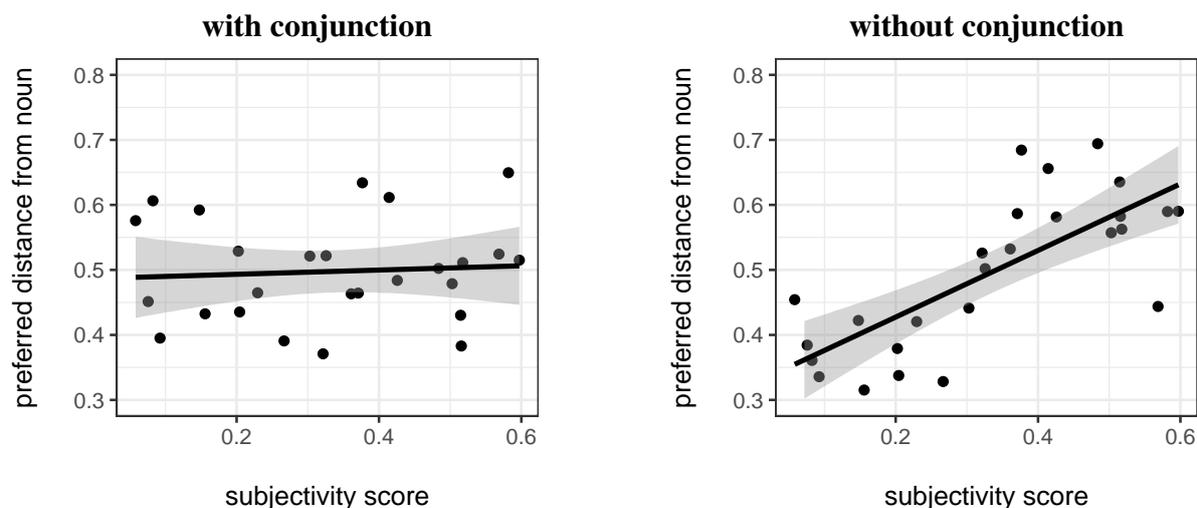


Figure 7: Spanish ordering preferences plotted against subjectivity scores for each of the 26 adjectives tested. *Left*: ordering preferences with conjunction; subjectivity accounts for 1% of the variance in the ordering preferences ($r^2 = 0.01$, 95% CI = [0.00, 0.06]). *Right*: ordering preferences without conjunction; subjectivity accounts for 56% of the variance in the ordering preferences ($r^2 = 0.56$, 95% CI = [0.23, 0.73]).

above, with Spanish translations of the instructions and test items. For each of the 26 adjectives tested, we computed a mean faultless disagreement score, which we use below to compare with the ordering preferences that were measured.

Comparing ordering preferences with subjectivity. We have two cases to check: first, ordering preferences with conjunction, and second, ordering preferences without conjunction. With conjunction, subjectivity fails to predict ordering preferences (Figure 7, *left*). This prediction failure arises because, as we saw above, there are no ordering preferences to predict in Spanish when multi-adjective strings are formed via conjunction. In contrast, without conjunction, subjectivity is a robust predictor of ordering preferences (Figure 7, *right*). We therefore find the prediction regarding conjunction borne out: conjunction neutralizes ordering preferences, at least in Spanish. Without conjunction, we find subjectivity-based ordering preferences post-nominally, a finding also in line with the predictions of incremental semantic restriction as a driver of subjectivity-based preferences.

3.3. Arabic

While Spanish stands apart with its preference for conjunction in multi-adjective strings, it also stands apart—at least relative to English and Tagalog—with its post-nominal adjectives. Here we test another language with post-nominal adjectives: Arabic.¹⁰ Unlike Spanish, Arabic speakers do not report a general preference for conjunction in multi-adjective strings.

¹⁰The material in this section summarizes portions of the findings from Kachakeche and Scontras (2020).

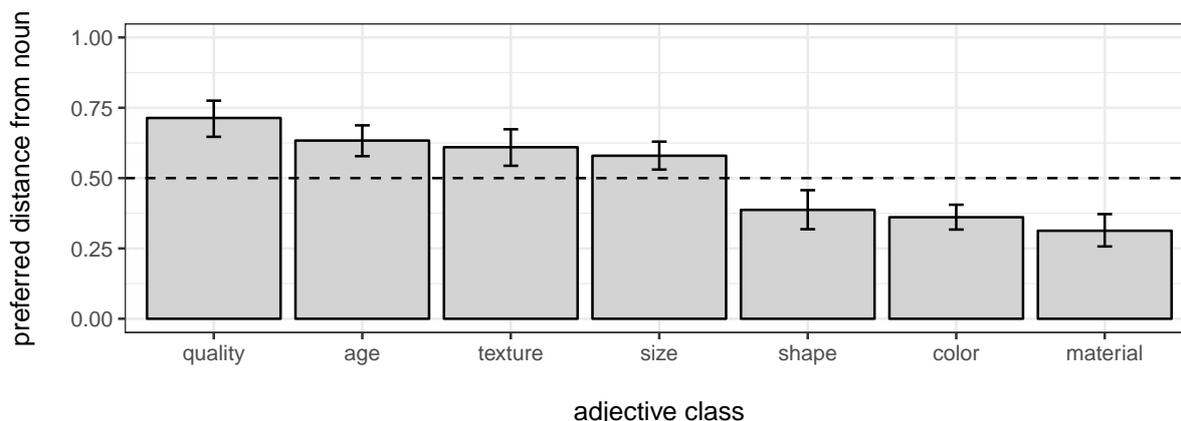


Figure 8: Arabic naturalness ratings grouped by adjective semantic class.

3.3.1. Measuring preferences

We replicated the ordering preferences experiment (without conjunction) using Arabic translations of the English materials.¹¹ Given that participants indicated preferences for written strings, testing proceeded using Modern Standard Arabic.

Participants. We recruited 135 participants through Mechanical Turk; 24 were identified as native speakers of Arabic on the basis of a post-test questionnaire.

Procedure. The procedure was identical to the ordering preferences experiments presented above, with the exception that instructions and materials appeared in Arabic.

Results. Figure 8 plots average preferred-distance measures, demonstrating clear evidence of stable ordering preferences in Arabic: some classes of adjectives are preferred farther from the noun, while others are preferred closer.

3.3.2. Measuring subjectivity

To measure subjectivity, we replicated the faultless disagreement task using Arabic materials.

Participants. We recruited 135 participants through Mechanical Turk; 16 were identified as native speakers of Arabic.

Procedure. The procedure was identical to the faultless disagreement experiments described above, with Arabic translations of the instructions and test items. For each adjective, we computed a mean faultless disagreement score by averaging participants' responses.

Comparing ordering preferences with subjectivity. Figure 9 plots adjective ordering preferences against the mean subjectivity scores for each of the adjectives tested. In Arabic, subjectivity is a robust predictor of adjective ordering preferences, thereby further confirming the prediction that subjectivity-based ordering preferences should arise whenever the hierarchical structure of modification supports them, whether pre- or post-nominally.

¹¹See Kachakeche and Scontras (2020) for details.

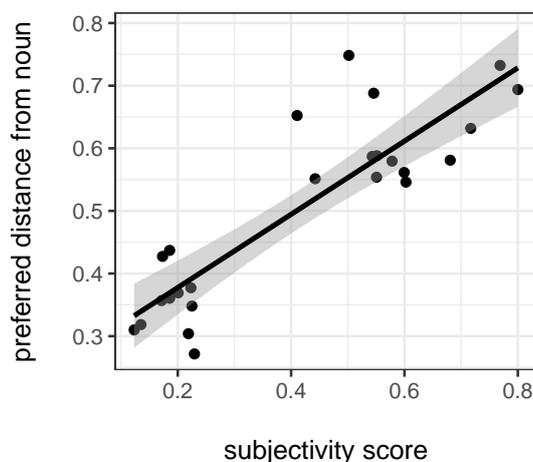


Figure 9: Arabic ordering preferences plotted against subjectivity scores for each of the 25 adjectives tested. Subjectivity accounts for 76% of the variance in the ordering preferences ($r^2 = 0.76$, 95% CI = [0.57, 0.88]).

3.4. English conjunction

So far, we have found clear evidence of stable ordering preferences despite the obligatory presence of LINKER in Tagalog, stable ordering preferences in post-nominal strings in both Spanish and Arabic, and no ordering preferences in the presence of conjunction in Spanish. With the results concerning the absence of stable preferences with conjunction in Spanish, we set out to confirm reports in the literature that conjunction neutralizes ordering preferences also in English (Ford and Olson, 1975; Byrne, 1979).

3.4.1. Measuring preferences

We re-ran the English ordering preferences experiment from Scontras et al. (2017), this time conjoining adjectives in multi-adjective strings.

Participants. We recruited 50 participants through Mechanical Turk; 49 were identified as native speakers of English.

Procedure. The experiment was a direct replication of the ordering preferences experiment from Scontras et al., with the exception that participants rated multi-adjective strings formed with conjunction.

Results. Figure 10 plots average preferred-distance measures. Unlike in Spanish, English speakers continue to have stable ordering preferences even with conjunction. Moreover, the English conjunction ratings closely replicate the qualitative results from the conjunction-free baseline observed by Scontras et al.

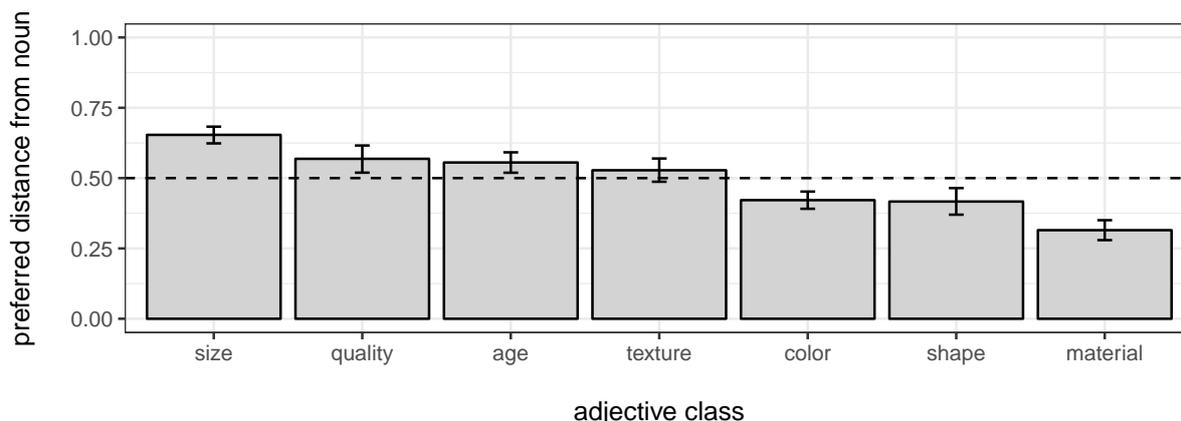


Figure 10: English naturalness ratings for conjoined multi-adjective strings grouped by adjective semantic class.

3.4.2. Comparing ordering preferences with subjectivity as measured by Scontras et al. (2017)

Having documented stable ordering preferences with conjunction in English, next we check the extent to which subjectivity predicts those preferences. To do so, we used the subjectivity scores measured by Scontras et al.; as Figure 11 shows, subjectivity continues to be a robust predictor of ordering preferences with conjunction in English ($r^2 = 0.68$; 95% CI [0.45, 0.80]). This finding stands at odds with our finding that conjunction neutralizes ordering preferences in Spanish, as well as with previous claims about the role of conjunction in English ordering preferences (Ford and Olson, 1975; Byrne, 1979). However, it is true that ordering preferences are less robust with conjunction than without it: as Figure 11 shows, while subjectivity accounts for 68% of the variance in the conjunction preferences, it accounts for 85% of the variance in preferences without conjunction, in large part because the preferences without conjunction are more extreme.

4. Discussion

The explanation for subjectivity-based adjective ordering preferences in terms of incremental semantic restriction (Simonič, 2018; Franke et al., 2019; Scontras et al., 2019) makes clear predictions about where we should expect to find communicative pressure toward these preferences. First, whenever adjectives incrementally restrict a nominal denotation—whether pre- or post-nominally—there should be pressure toward subjectivity-based ordering. Second, in the absence of incremental semantic restriction—as in the case of conjoined adjectives—such pressures should not obtain. In this paper, we set out to test both sets of predictions using cross-linguistic data obtained in a series of eight experiments.

First, regarding where we should expect to find preferences, we tested pre-nominal adjectives in Tagalog formed with the obligatory LINKER. Even if LINKER receives a semantics similar to conjunction, as some have proposed (Rubin, 1994; Scontras and Nicolae, 2014), its structure is such that adjectives compose incrementally with the nominal: adjectives closer to the modified noun restrict the nominal denotation before adjectives that are farther away. Thus, we should expect that speakers prefer to use less error-prone, less subjective adjectives earlier in this incremental process, and, indeed, we found just this pattern: in Tagalog, speakers have robust

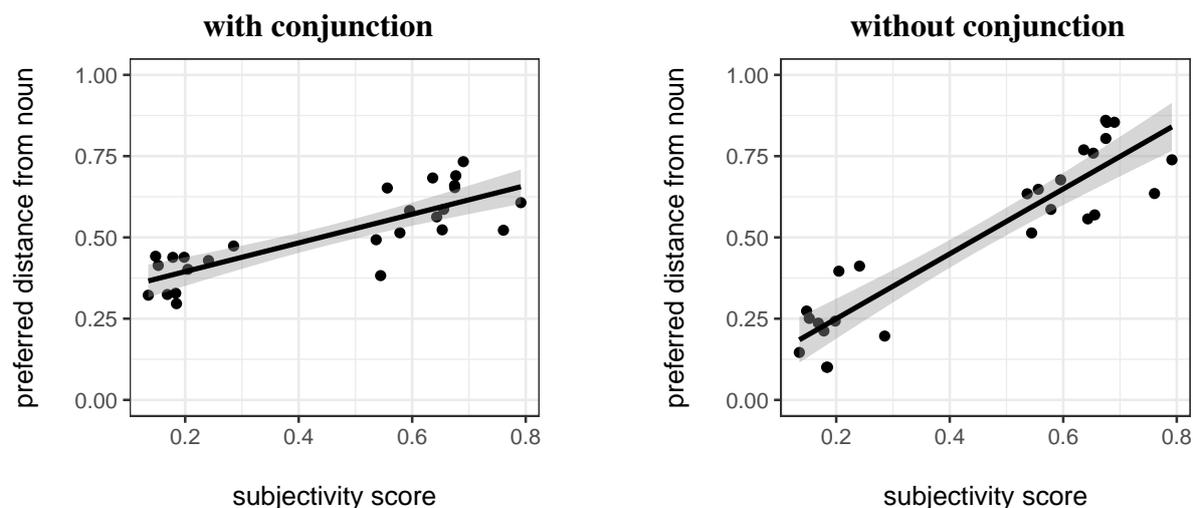


Figure 11: English ordering preferences plotted against subjectivity scores for each of the 26 adjectives tested. *Left*: ordering preferences with conjunction; subjectivity accounts for 68% of the variance in the ordering preferences ($r^2 = 0.68$, 95% CI = [0.45, 0.80]). *Right*: ordering preferences without conjunction from Scontras et al. (2017); subjectivity accounts for 85% of the variance in the ordering preferences ($r^2 = 0.85$, 95% CI = [0.75, 0.90]).

subjectivity-based adjective ordering preferences.

We should also expect to find subjectivity-based adjective ordering preferences in languages with post-nominal adjectives where linear distance tracks the hierarchical structure of semantic composition, for example in Spanish and Arabic. We found straightforward evidence of post-nominal subjectivity-based ordering preferences in Arabic. We also found such preferences in Spanish. However, in Spanish, the empirical picture is complicated by a preference for conjunction in multi-adjective strings. Only in strings formed without conjunction do we find stable ordering preferences.

The point about conjunction confirms our second prediction regarding where we should not expect to find subjectivity-based ordering preferences. In Spanish, conjunction neutralizes ordering preferences. This effect makes sense if the pressure for ordering preferences comes from a desire to compose less subjective adjectives earlier with the modified noun; with conjunction, as in (4), the adjectives make their semantic contribution together after they are conjoined, so pressures mediating the order in which adjectives compose cannot apply. (It is not clear how one would explain this result under a memory-based ordering account, given that relative distance is preserved with conjunction; *pace* Hahn et al., 2018.)

But the story on conjunction does not end there. In English, preferences weaken but persist in the presence of conjunction. One way to understand the English result is that in languages where multi-adjective strings optionally feature conjunction (as in English), the regularity introduced in conjunction-less strings can bleed over to strings with conjunction. English speakers thus internalize the statistical ordering regularity from non-conjoined adjective strings and use that knowledge to inform preferences for conjoined strings. In Spanish, where speakers commonly prefer conjunction in multi-adjective strings, there may be less of a source for an ordering regularity—conjunction-free strings are dispreferred and therefore occur less frequently—that could be extended by analogy to the conjoined strings.

Before concluding, there is one more point that warrants comment. While we have focused on the similarities across languages in terms of subjectivity-based ordering preferences, there are also differences in our results. Some languages exhibit a stronger correlation between their ordering preferences and adjective subjectivity (e.g., Arabic, English), while others exhibit weaker correlations (e.g., Tagalog, Spanish). At the level of adjective classes, we find qualitative differences in ordering preferences across languages. While these differences no doubt arise in part because of the idiosyncratic details of our studies (e.g., noise in our data introduced by the availability of subjects), there are likely other, more meaningful factors at play. For example, if we follow Franke et al. (2019) in assuming that ordering preferences develop as speakers internalize the statistical regularities of their input, if the input contains fewer multi-adjective strings, there will be less data to serve in the formation of preferences, and so we might predict weaker preferences overall—a possibility in the case of the weaker preferences observed in conjunction-free strings in Spanish. We leave it to future research to more fully explore this issue.

5. Conclusion

Our results provide further support for the empirical generalization concerning subjectivity in adjective-ordering preferences, as well as support for the role of incremental semantic restriction in subjectivity-based ordering preferences. In Tagalog and Arabic, where adjectives incrementally restrict the nominal denotation, we find subjectivity-based preferences regardless of whether adjectives appear with linking particles or whether they appear pre- or post-nominally. In Spanish, the presence of subjectivity-based preferences depends on whether multi-adjective strings are formed via conjunction: only without conjunction do we find stable ordering preferences in Spanish, and those preferences track adjective subjectivity. Thus, when conjunction disrupts the hierarchical structure that would deliver incremental restriction, pressure toward subjectivity-based ordering disappears. The complicating factor is the picture from English, where preferences weaken but persist with conjunction, suggesting that the regularity introduced in conjunction-less strings can bleed over to strings with conjunction if there exists sufficient support in the input.

References

- Bever, T. G. (1970). The cognitive basis for linguistic structures. In R. Hayes (Ed.), *Cognition and Language Development*, pp. 279–362. New York: Wiley & Sons, Inc.
- Byrne, B. (1979). Rules of prenominal adjective order and the interpretation of “incompatible” adjective pairs. *Journal of Verbal Learning and Verbal Behavior* 18(1), 73–78.
- Cinque, G. (1994). On the evidence for partial N-movement in the Romance DP. In R. S. Kayne, G. Cinque, J. Koster, J.-Y. Pollock, L. Rizzi, and R. Zanuttini (Eds.), *Paths Towards Universal Grammar. Studies in Honor of Richard S. Kayne*, pp. 85–110. Washington DC: Georgetown University Press.
- Dixon, R. (1982). *Where have all the adjectives gone? And other essays in semantics and syntax*. Berlin: Mouton.
- Foley, W. (1975). *Comparative Syntax in Austronesian*. Ph. D. thesis, University of California, Berkeley.
- Ford, W. and D. Olson (1975). The elaboration of the noun phrase in children’s description of objects. *Journal of Experimental Child Psychology* 19(3), 371–382.

- Franke, M., G. Scontras, and M. Simonič (2019). Subjectivity-based adjective ordering maximizes communicative success. In *Proceedings of the 41st annual meeting of the Cognitive Science Society*, pp. 344–350.
- Hahn, M., J. Degen, N. D. Goodman, D. Jurafsky, and R. Futrell (2018). An information-theoretic explanation of adjective ordering preferences. In *Proceedings of the 40th annual meeting of the Cognitive Science Society*, pp. 1766–1771.
- Hetzron, R. (1978). On the relative order of adjectives. In H. Sella (Ed.), *Language Universals*, pp. 165–184. Tübingen: Narr.
- Hill, F. (2012). Beauty Before Age? Applying Subjectivity to Automatic English Adjective Ordering. In *Proceedings of the NAACL HLT 2012 Student Research Workshop*, pp. 11–16.
- Kachakeche, Z. and G. Scontras (2020). Adjective ordering in Arabic: Post-nominal structure and subjectivity-based preferences. *Proceedings of the Linguistic Society of America* 5, 419–430.
- Kölbel, M. (2004). Faultless Disagreement. *Proceedings of the Aristotelian Society* 104, 53–73.
- MacFarlane, J. (2014). *Assessment Sensitivity*. Oxford: Clarendon Press.
- Martin, J. E. (1969). Some competence-process relationships in noun phrases with prenominal and postnominal adjectives. *Journal of Verbal Learning and Verbal Behavior* 8, 471–480.
- Rosales Jr., C. M. and G. Scontras (2019). On the role of conjunction in adjective ordering preferences. *Proceedings of the Linguistic Society of America* 4(32), 1–12.
- Rubin, E. J. (1994). *Modification: A Syntactic Analysis and Its Consequences*. Ph. D. thesis, Cornell University, Ithaca, NY.
- Samonte, S. and G. Scontras (2019). Adjective ordering in Tagalog: A cross-linguistic comparison of subjectivity-based preferences. *Proceedings of the Linguistic Society of America* 4(33), 1–13.
- Schmidt, L. A., N. D. Goodman, D. Barner, and J. B. Tenenbaum (2009). How tall is *tall*? Compositionality, statistics, and gradable adjectives. In *Proceedings of the 31st annual meeting of the Cognitive Science Society*, pp. 3151–3156.
- Scontras, G., J. Degen, and N. D. Goodman (2017). Subjectivity predicts adjective ordering preferences. *Open Mind: Discoveries in Cognitive Science* 1(1), 53–65.
- Scontras, G., J. Degen, and N. D. Goodman (2019). On the grammatical source of adjective ordering preferences. *Semantics and Pragmatics* 12(7).
- Scontras, G. and A. C. Nicolae (2014). Saturating syntax: Linkers and modification in Tagalog. *Lingua* 149, 17–33.
- Scott, G.-J. (2002). Stacked adjectival modification and the structure of nominal phrases. In G. Cinque (Ed.), *The cartography of syntactic structures, Volume 1: Functional structure in the DP and IP*, pp. 91–120. Oxford: Oxford University Press.
- Simonič, M. (2018). Functional explanation of adjective ordering preferences using probabilistic programming. Master's thesis, University of Tübingen.
- Sproat, R. and C. Shih (1991). The cross-linguistic distribution of adjective ordering restrictions. In C. Georgopoulos and R. Ishihara (Eds.), *Interdisciplinary approaches to language: Essays in honor of S.-Y. Kuroda*, pp. 565–593. Dordrecht: Kluwer Academic Publishers.
- Tucker, G. (1998). *The Lexicogrammar of Adjectives: A Systemic Functional Approach to Lexis*. London: Cassell Academic.
- Whorf, B. L. (1945). Grammatical Categories. *Language* 21(1), 1–11.