Phonetic effects of grammatical category: How category-specific prosodic phrasing and lexical frequency impact the duration of nouns and verbs

Arne Lohmann^{a*} and Erin Conwell^b

*Corresponding author's contact info:

Arne Lohmann, Department of English and American Studies, Heinrich-Heine-Universität Düsseldorf, Universitätsstrasse 1, 40225 Düsseldorf, Germany, E-mail: arne.lohmann@hhu.de, Tel: +49-211-81-15970

^aDepartment of English and American Studies, Heinrich-Heine-Universität Düsseldorf, Düsseldorf, Germany

^bDepartment of Psychology, North Dakota State University, Fargo, USA

Manuscript Accepted for Publication at *Journal of Phonetics*:

<https://doi.org/10.1016/j.wocn.2019.100939>

Abstract

This paper is concerned with phonetic correlates of grammatical category, specifically the finding that nouns are pronounced with greater duration than verbs in discourse. Most previous research has attributed this difference to the sentence positions that the two grammatical categories occupy and concomitant prosodic effects. Based on previous findings, we test two further effects, namely a category-specific effect on prosodic phrasing, which leads to stronger prosodic boundaries after nouns than verbs even in maximally similar syntactic contexts, and a reductive effect of lexical frequency leading to shorter durations of the more frequent word. These effects are tested in a production study investigating durational differences of twelve noun-verb homophone pairs in English in two clause-medial contexts. We find evidence for both effects: prosodic boundaries are stronger after nouns than verbs across all conditions, resulting in greater durations of nouns due to pre-boundary lengthening. Furthermore, differences in frequency result in a reduced duration of the homophone of the pair which has the greater frequency. We propose an explanation in which phonetic effects of grammatical category are caused by the interplay of sentence prosody, category-specific prosodic phrasing and lexical frequency.

Keywords

Grammatical categories, prosody, phrasing, acoustic phonetics, frequency effects, speech production, homophones

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

Phonetic effects of grammatical category: How category-specific prosodic phrasing and lexical frequency impact the duration of nouns and verbs

1. Introduction

The present study investigates how the grammatical category of a word affects its acoustic realization. Specifically, we ask what underlies the greater duration of nouns as compared to verbs in discourse, the acoustic contrast that constitutes the most conspicuous and robustly evidenced phonetic difference between the two categories (Coker, Umeda, & Browman, 1973; Conwell & Morgan, 2012; Dिल्ts, 2013; Lightfoot, 1970; Sorensen, Cooper, & Paccia, 1978; Yao, 2011).

Although differences in duration between nouns and verbs have been widely observed, the question of what is driving these differences has not been conclusively answered. In early work on the topic, Sorensen et al. (1978) and Cooper & Paccia-Cooper (1980) raise the question of whether grammatical category directly impacts phonetic realization, or whether durational differences are an epiphenomenon of other processes correlated with grammatical category. They argue for the latter and propose that the typical sentence positions in which nouns and verbs occur are the sole factor driving pronunciation differences. Nouns occur more frequently than verbs in constituent-final positions (i.e., at the end of phrases and clauses), while verbs typically occur in phrase-initial position (of VPs). These differences in position trigger two prosodic processes, both leading to longer duration of nouns than verbs. The first is the likelihood of attracting nuclear sentence stress or accent, as constituent-final position is the primary locus of sentence stress (as implemented in, e.g., the classic “nuclear stress rule” by Chomsky & Halle, 1968).¹ One acoustic correlate of accent is an increased duration via ‘accentual lengthening’ (Turk & White, 1999). The second relevant prosodic process is constituent-final lengthening, resulting in a greater duration of words preceding a syntactic/prosodic boundary (e.g., Wightman, Shattuck-Hufnagel, Ostendorf & Price, 1992). Due to their preferred sentence positions, nouns are more frequently affected by both of these processes, leading to longer duration as compared to verbs (see also the discussions in Cooper & Paccia-Cooper, 1980, and Sorensen et al., 1978). Sorensen et al. (1978, Experiment 2) elicited pronunciations of nouns and verbs in matched clause-final position, thus controlling for differences in accentuation and boundary-induced lengthening. They found that duration differences disappear, providing the crucial argument for their conclusion that pronunciation differences are solely due to sentence position (see also Cooper & Paccia-Cooper, 1980).

However, recent research indicates that the pronunciation of nouns and verbs may be influenced by further factors. Watson, Breen, & Gibson (2006) find that when nouns and verbs occur in syntactically similar clause-medial contexts, the likelihood of an intonational phrase boundary is greater after nouns than verbs. This suggests that there may be an independent effect on prosodic phrasing that relates to grammatical category, which leads to more pronounced prosodic boundaries after nouns than verbs and consequently longer durations of nouns even in matched syntactic contexts.

A second factor that may impact the acoustic realization of nouns and verbs is that there are differences in usage frequency between the two grammatical categories. In naturally occurring discourse, verbs have a lower type but higher token frequency than nouns. This means that the number of distinct noun types (i.e., distinct noun lemmas) is higher than the number of distinct verbs (i.e., distinct verb lemmas). However, each distinct noun has a lower frequency of occurrence (i.e., token frequency). To illustrate that point, in the COCA corpus (Davies, 2014),

¹ See Ladd (2008), Chapter 7 for a detailed overview and discussion of accounts of sentence accent.

which contains about 560 million words of text, the average noun lemma has a token frequency of 11,614 occurrences, while the average verb lemma occurs 28,232 times. Since it has been shown that lexical frequency yields a reductive effect on word duration (e.g., Gahl, 2008), one may expect shorter durations of verbs as compared to nouns. Wang & Xu (2017) argue that this factor is primarily responsible for acoustic differences between nouns and verbs, based on an analysis of productions of Mandarin noun-verb homophones. Evidence for lexical frequency co-determining the duration of noun-verb homophones in English has been found in child-directed corpora (Conwell, 2017) and also adult-directed speech (Lohmann, 2018).

The picture that emerges from these more recent findings is that the answer to the question of what underlies durational differences between nouns and verbs may be more complex than hitherto thought. It may be influenced by (1) their typical sentence positions and concomitant differences in accentual and pre-boundary lengthening, (2) a separate category-specific effect leading to differences in the strength of the following prosodic boundary, and (3) word frequency effects at the level of the individual category-specific word. The empirical challenge is that in naturally occurring speech all three of these factors may be jointly active and all of them predict a net result of shorter durations of verbs as compared to nouns. This makes it difficult to assess whether each of them independently impacts acoustic realization. The aim of the current study is to empirically test for independence of these factors by disentangling their effects. More specifically, we aim to put both the category-specific effect on prosodic phrasing as well as the reductive effect of lexical frequency (2 and 3 above) to a more rigorous test. This aim will be pursued by controlling for sentence position while testing for the individual contributions of prosodic boundary strength and frequency-induced reduction to word duration. In the following subsections we detail how these effects are hypothesized to impact the durational characteristics of nouns and verbs. We then test these effects in a production study of noun-verb homophones.

1.1 Nouns and verbs, prosodic boundary placement and pre-boundary lengthening

As discussed above, the different sentence positions of nouns and verbs lead to differences in duration, mostly due to a difference in degree to which the two grammatical categories are affected by pre-boundary lengthening. Nouns occur more frequently than verbs before syntactic boundaries of greater strength, which often coincide with prosodic boundaries (see e.g., Sorensen et al. 1978). The effect of these positional differences is that nouns are more frequently and more strongly affected by pre-boundary lengthening, as the degree of lengthening is positively correlated with the strength of the following prosodic boundary (Wightman, Shattuck-Hufnagel, Ostendorf, & Price, 1992).

One aim of the present paper is to test whether the distinction in grammatical category per se influences the prosodic phrasing of sentences even in controlled syntactic contexts, resulting in a greater lengthening of nouns. This idea builds on initial findings by Watson et al. (2006), who show that intonational and intermediate phrase boundaries are more likely after nouns than verbs even in well-controlled clause-medial contexts. In the following we explain the contribution of their research to the literature on the relationship between syntactic constituent structure and prosodic boundary strength, as well as how the present study expands on those findings.

A number of accounts have been developed to predict the size of the prosodic boundary between any two words within a sentence (see Ferreira, 2007). While the literature on the syntax-prosody relationship does not assume that syntactic and prosodic structure are isomorphic (see Selkirk, 2011; Shattuck-Hufnagel & Turk, 1996; Turk & Shattuck-Hufnagel, 2014), all of these accounts agree that the syntactic structure of the sentence determines prosodic phrasing to a considerable degree.

An early approach is Selkirk's (1984) model of “silent demibeat additions” (SDA), which are silent positions that are added after syntactic constituents, the number of which indicates the strength of the prosodic boundary. The number of silent beats is dependent on the distinction between function and content words (but not between nouns and verbs) and the position within the sentence and phrase. For instance, more silent beats are added after phrase-final than phrase-medial words and at the right edge of phrases that are daughter phrases of the sentence node than after weaker syntactic boundaries (see Selkirk, 1984, pp. 313-319 for operationalization details). One weakness of the SDA approach is that it only considers linguistic material preceding the target word, but not upcoming material, although the size and complexity of upcoming constituents also influences prosodic phrasing (see e.g., Ferreira, 2007, see also D’Imperio, Giorka, Frota, Prieto, & Vigário, 2005; Frota, D’Imperio, Elordieta, Prieto, & Vigário, 2010, on additional language-specific effects).

Other accounts have been developed that remedy this shortcoming. These accounts rely on algorithms that output a certain value indicating the break or boundary strength at an inter-word location. The most well-known of these are the algorithms by Gee & Grosjean (1983) and Watson & Gibson (2004, 2005). Both are based on grouping the sentence into phonological phrases, which are defined as “all the words within the maximal projection of a lexical head on the lexical head’s left side” (Watson & Gibson, 2004, pp.730-731). The “left/right boundary” (LRB) algorithm by Watson & Gibson (2004, 2005) takes into account the number of phonological phrases to the left and to the right of a given location. The LRB strength is calculated for each boundary between two phonological phrases and is based on the complexity of the syntactic constituent ending at that point and the complexity of the upcoming constituent, measured in number of phonological phrases. Phonological phrases in the upcoming constituent that are arguments of the preceding word are not counted. An example application of their algorithm is exemplified below (copied from Watson & Gibson, 2004, p. 734, see also Ferreira, 2007):

(1) (The judge)₄(who the reporter)₂(for the newspaper)₄ (ignored)₇(fired)₁(the secretary).

In example (1) above, the strongest prosodic boundary is predicted to occur between the subject NP and the VP, which is also the strongest syntactic boundary within the sentence. While the LRB and other algorithms have been shown to fairly successfully predict prosodic boundary strength, it is clear that syntactic constituent structure is not the only factor impacting prosodic phrasing. A further factor that is implemented in the LRB algorithm is whether an argument or adjunct follows a word, with more pronounced prosodic boundaries preceding the latter than the former. Further exploring the role of argumenthood, Watson et al. (2006) show that both semantic closeness as well as syntactic obligatoriness of an upcoming constituent impact prosodic boundary strength. Crucially, in the same paper, Watson et al. (2006) obtain initial evidence that grammatical category also affects intonational phrasing, with prosodic boundaries having greater strength after nouns than verbs, all other things being equal.

In a subset of their test sentences, Watson et al. (2006) controlled for characteristics such as argument/adjuncthood and obligatoriness of the following constituent across grammatical category and measured prosodic boundary strength after the target word employing the ToBI framework (Beckman & Ayers Elam, 1997; Beckman & Hirschberg, 1994). An example pair of their test sentences is the following. Here the target words *arrival/arrived* are followed by an optional adjunct PP.

(2a) The reporter’s *arrival* after the crash unnerved the officials. (N condition)

(2b) The reporter *arrived* after the crash and this unnerved the officials. (V condition)

Applying the LRB algorithm to these sentences yields the same predicted boundary strength of 2 for both the noun and the verb condition, as in both cases the syntactic constituent (NP or VP,

respectively) continues after the target word with a right-hand constituent that consists of one phonological phrase (the PP), which contributes a break value of 1 to the location. Moreover, the location constitutes the end of a phonological phrase, which adds a further break value. When coding the prosodic boundary strength in data obtained through a reading study, Watson et al. (2006) find a considerable degree of variation in the prosodic boundaries following the target words. Crucially, however, they report a probabilistic effect of greater prosodic boundary strength after the noun as compared to the verb. Although the acoustic realization of the target words was not investigated by Watson et al. (2006), such an effect is likely to result in greater durations of nouns via pre-boundary lengthening.

One aim of the present study is to replicate the effect found by Watson et al. (2006) and explore its explanatory value for durational differences between nouns and verbs. To that end we constructed test sentences similar to the ones employed by Watson et al. (2006), but containing noun-verb homophones. An exemplary sentence pair is given in (3), below.²

In these test sentences the noun and the verb condition are identical regarding the sequence of segments until the end of the PP that follows the target words. Note, however, that the syntactic constituency structure differs between the two conditions. The target nouns are part of an –s genitive NP, while the target verb introduces a VP following a subject NP (compare 3 a, b). As noted above, Watson & Gibson's (2004) LRB algorithm still outputs the same predicted boundary strength for both conditions. However, this result holds only when phonological phrases are taken as the unit of measurement for constituent complexity. Watson & Gibson acknowledge that other ways of measuring complexity may also work and could even be better suited (see Watson & Gibson, 2004, p. 731). Alternative units include the number of words or syllables making up the crucial syntactic constituent (see e.g., D'Imperio et al., 2005). When counting words or syllables, the noun and the verb conditions of the test sentences differ in the leftward complexity of the constituent (NP or VP) that contains the target words. While the verb is the initial word of the syntactic constituent (i.e., the VP) that it is part of, this is not the case in the noun condition, in

² The entire set of test sentences employed by Watson et al. (2006) were designed to include differences with regard to argumenthood and obligatoriness of the following syntactic material. It is not clear whether nouns exhibit the same tight relationship to arguments as verbs (see Watson et al. 2006, p. 1047), a difference that may co-vary with the category effect when testing sentences in which the target words are followed by arguments. To focus only on possible effects of grammatical category, we held these variables constant. In our test sentences the target words are always followed by an optional, non-argument PP.

which the possessor noun and a determiner precede the target word, resulting in greater leftward complexity of the embedding syntactic constituent (i.e., the NP). Another possibly relevant difference in syntactic constituent structure is that the length/complexity of the clauses that the target words are part of differs between the noun and the verb condition. In the noun condition, clause and sentence are congruent; however, in the verb condition the embedding clause is shorter as it ends after the PP, before the conjunction *and* (cf. 3a vs. 3b). Because constituents of greater size are more likely to be split up by prosodic boundaries of greater strength (Turk 2010: 243-245), this difference may lead to a boundary of greater strength after the noun as opposed to the verb. In sum, it remains possible that the finding of a greater prosodic boundary strength after nouns as compared to verbs is an artefact of the difference in constituent structure between the two conditions and not an effect of grammatical category per se.

Due to this shortcoming of the test sentences employed by Watson et al. (2006), we tested the effect of grammatical category on prosodic phrasing in an additional syntactic environment in which the constituent structure was more carefully matched between the two conditions. Similar to the examples given in (3) the target words occurred in clause-medial position preceding a non-argument, non-obligatory prepositional phrase. We kept the lexical material as similar as possible across the two categories by using lexically identical prepositional phrases and by also maximizing the similarity of the lexical material preceding the target words. Crucially, the syntactic structure to the left of the target words was also kept maximally similar across category, which meant that the complexity of the embedding constituents was kept constant with regard to the number of syllables. See the examples in (4a, b) for illustration.

Half of the target sentences are like the examples (4 a, b), in that they differ only in the function word introducing the phrasal constituent containing the target word, which starts either with the determiner *a*, or the infinitive marker *to*. The other half of the sentence pairs were similarly matched with regard to constituent complexity, but required further lexical modifications to arrive at natural contexts for both the noun and the verb (see for example 5a, b).

- (5a) I kindly suggest you read this *book* for the performance.
 (5b) This is the ticket you want to *book* for the performance.

If the finding by Watson et al. (2006) was solely due to the difference in constituent structure of the test sentences, we would expect to find differences in prosodic boundary strength and duration of the target words only in that context, but not in the more carefully controlled context specified in examples (4) and (5). If both the constituent complexity of the embedding constituent and grammatical category yield independent effects, we would expect to find differences in prosodic boundary strength between the two categories in both environments with a stronger effect in the test sentences modeled after Watson et al. (2006).

1.2 Nouns and verbs and reductive effects of lexical frequency

The second aim of the present study is to test the effect of lexical frequency on the duration of noun-verb homophones. Evidence for a reductive effect of frequency on word duration in general has been provided in a number of studies (see Jurafsky, 2003). Whether the same effect exists within homophone pairs, has not been clearly established. Gahl (2008) reports a reductive effect differentiating the durations of homophonous words of different frequency (e.g., *thyme* vs. *time*) and similar evidence for an effect of category-specific frequency has been obtained for English noun-verb homophones (Conwell, 2017a; Lohmann, 2018). However, these studies relied on corpora of spontaneous speech, which limits the validity of the results, especially in case of noun-verb homophones. One shortcoming of corpus data is that the sentence positions of the target words cannot be held constant, but rather systematically co-vary between nouns and verbs, due to the different occurrence patterns discussed above. Secondly, lexical frequency and grammatical category are correlated in natural discourse (see Section 1). Both these correlations are avoided in the present study. This is achieved by employing both noun-dominant homophone pairs (i.e., the noun is more frequent than the verb), as well as verb-dominant homophone pairs and by analysing productions of noun-verb homophones in maximally similar syntactic contexts. As a result, the frequency hypothesis is put to a much more rigorous test.

2. Materials and methods

2.1 Materials and design

2.1.1 Target words

We chose target noun-verb pairs based on the following criteria. Because we wanted all words to be homogenous in terms of word shape, we chose to include only monosyllabic words with a filled onset and coda. To create a dataset with maximally different vowels, we selected words featuring the three vowels /æ/, /ɪ/ and /ʊ/.³ To test the effect of category-specific frequency, we included noun-verb homophones that cover the spectrum from very noun-dominant to very verb-dominant. Finally, word pairs were selected for their ability to occur naturally in both grammatical categories in the two syntactic contexts employed (see Sections 1.1 and 2.1.2). These criteria led to the twelve target word pairs in Table 1. There are four target word pairs for each vowel.

³ Because most of our participants are from North Dakota, a region of the US which is difficult to categorize with regard to the *cot/caught* merger (i.e., some speakers exhibit it and some do not), we did not include low back vowels.

Table 1 *The target word pairs ordered by frequency bias (verb-dominant to noun-dominant) and vowel phoneme*

N/V pair	<i>push</i>	<i>look</i>	<i>dig</i>	<i>kick</i>	<i>kiss</i>	<i>pack</i>	<i>cook</i>	<i>chat</i>	<i>tip</i>	<i>nap</i>	<i>snack</i>	<i>book</i>
Logged noun ratio	-1.26	-1.10	-0.74	-0.65	-0.34	-0.25	-0.24	0.12	0.28	0.84	1.06	1.19
N/V pair	<i>push</i>	<i>look</i>	<i>cook</i>	<i>book</i>	<i>pack</i>	<i>chat</i>	<i>nap</i>	<i>snack</i>	<i>dig</i>	<i>kick</i>	<i>kiss</i>	<i>tip</i>
Vowel phoneme			ʊ			æ					ɪ	

The logged noun ratio given in Table 1 is the log-transformed (base 10) ratio of the noun frequency divided by the verb frequency. A negative value of -1 indicates that the verb is ten times as frequent as the noun, while a positive value of 1 indicates that the noun is ten times as frequent as the verb. The calculation is based on the lemma frequency counts of the words in the SUBTLEX-US corpus, that is, the frequency of all inflectional variants of the noun and the verb, respectively (Brysbaert & New, 2009; Brysbaert, New, & Keuleers, 2012). SUBTLEX-US is a corpus of subtitles from US film and TV productions, mostly from the latter part of the 20th century. Subtitle frequencies may more accurately reflect everyday language exposure than frequencies obtained from more literary sources (see the discussion in Brysbaert & New, 2009).

2.1.2 Construction and presentation of target sentences

As described above, we employed two different syntactic environments in which the target words occurred: one modelled after Watson et al. (2006) and one in which syntactic constituent structure was still more carefully matched between the noun and the verb condition. Target sentences for each word pair and type of environment were designed as explained in Section 1.1. An additional factor that we manipulated was the length and complexity of the PP following the target words by optionally adding an additional PP that was syntactically integrated into the first one (see the material in parentheses below). This was done to enhance the variability of the prosodic boundary following the target words, as the length of an upcoming constituent is positively correlated with the strength of the prosodic boundary preceding it (see Watson et al., 2006; Watson & Gibson, 2004; Turk, 2010). Consequently, we created a long and a short version of each sentence. Henceforth, we will refer to the more carefully controlled syntactic environment as ‘context A’, and to the syntactic environment modelled after Watson et al. (2006), as ‘context B’.

Context A:

- (6a) The kids began a chat in front of the museum (of Natural History).
- (6b) The kids began to chat in front of the museum (of Natural History).

Context B:

- (7a) The actors’ kiss on the movie set (for the new production) annoys the director.
- (7b) The actors kiss on the movie set (for the new production) and this annoys the director.

As discussed above, in context B (modelled after Watson et al. 2006), the syntactic constituent structure of the sentence is not identical between the two categories. Apart from a difference in prosodic boundary strength after the target words (see Section 1.1), this may have further consequences for the overall prosodic structure of the target sentences. In designing the target sentences for context A, we tried to keep the global prosodic characteristics as similar as possible, in particular the occurrence of pitch accents in the vicinity of the target word and the target word itself. We expect the target word itself to be spoken with a pitch accent across all contexts and conditions in this study.

To further investigate possible differences in accentuation between the noun and the verb condition, we carried out a sample analysis of the target sentence productions by four speakers (2 male, 2 female; sampled from the total number of 78 participants, see section 2.2). We found the target words to be pronounced with a pitch accent across the two conditions in both syntactic contexts, with only three exceptions (2 verbs, 1 noun; out of 196 total target sentence productions from these four speakers). The PP following the target words was characterized by at least one further pitch accent, typically on the final noun, (e.g., *museum* or *History* in examples 6a and 6b). The analysis did not indicate a systematic difference with regard to pitch accent occurrence between the noun and verb condition.

While the target words are almost uniformly pitch-accented between the noun and the verb condition, there remains the possibility of a systematic difference in the type of pitch accent on the target word. Our sample analysis does not indicate such a difference. In most cases the pitch accent on the target word is a high tone (H* in ToBI notation) in both the noun and the verb condition (see also Figures 2 and 3 for examples of a pitch accent transcription of target sentences).

We also calculated the target words' F0 range for the individual conditions as a measure that may capture one aspect of the phonetic implementation of pitch accents. More pronounced pitch accents should be marked by a greater pitch excursion and consequently a greater F0 range. See Table 2, which displays the F0 range of the target words in semitones by grammatical category and syntactic environment (F0 range was operationalized as the difference between the F0 maximum and the F0 minimum at the word level.)

Table 2 F0 range by grammatical category and type of sentence context

	Context A:	<i>t</i>	<i>p</i>	Context B:	<i>t</i>	<i>p</i>
	Nouns / Verbs			Nouns / Verbs		
F0 range (semitones)	6.29 / 6.31	-0.06	>0.9	5.72 / 5.09	2.15	<0.05

In the context designed by us (context A), we observe almost identical F0 range values. Hence, F0 range does not indicate systematic differences in the phonetic realization of pitch accents between the noun and the verb condition. However, the result for context B indicates a small difference in F0 range, hinting at the possibility of a more pronounced pitch excursion of the accents that are placed on the target nouns in the context modelled after Watson et al. (2006). In order to statistically control for this difference, we use the target words' F0 range as a control variable in the models predicting duration (see Section 3.4; see also Section 4.4. on the issue of whether the type of pitch accent may affect the results).⁴

Regarding the presentation of the target sentences for the reading task, we embedded the target sentences in short paragraphs in order to avoid unnatural list reading. The target sentence was always the last sentence of the paragraph and was preceded by two sentences that provided a semantic context for the target sentence. For example, the paragraph for the pair *dig(v)/dig(n)* in context A was the following:

⁴ The editor suggested that we explore whether the difference in F0 range in context B is due to a systematic difference in the likelihood of narrow focus on the target word, marked by a contrastive focal accent. We investigated the 96 target sentences of context B produced by the four speakers in the sub-sample we analysed. None of the target sentences were perceived as being pronounced with contrastive accent on the target word. Of course, this leaves open the possibility of other differences in pitch accent occurrence between the two conditions that we did not explore (see also the discussion in Section 4.4).

- (8) A famous Spanish pirate was rumored to have buried his treasure near the ocean. After discovering an old map with a big X, Ricki thought she knew where to look. The treasure hunter began {the dig, to dig} behind the rock on the beachfront.

2.1.3 Creation of lists for reading

Each participant produced both a noun sentence and a verb sentence for each word pair in both types of syntactic context. To reduce priming or repetition effects, we ensured that no speaker produced both versions of a minimally different sentence pair (e.g., 6a and 6b, above). We created two lexically distinct, but structurally similar sentences for each word pair in each type of context. The number of syllables of the constituents preceding the target word and of the PP following it (both the short and extended versions) were kept constant across the two versions (e.g., 9a, b and 10 a, b; the complete stimulus list is in the Appendix).

- (9a) The kids began a chat in front of the museum (of Natural History).
(9b) Their friends started to chat right behind the old statue (of President Kennedy).
(10a) The parents' chat during the play (on US history) angers Mr. Robinson.
(10b) The students chat about the quiz (on advanced chemistry) and this makes them feel more confident.

We systematically varied the particular configuration of grammatical category and length of the following PP for each word pair across participants by creating four lists, each containing a total of 48 target paragraphs (four for each of the 12 target word pairs), a number which fits in with suggestions for best practice in conducting speech elicitation tasks (Niebuhr & Michaud, 2015). To create the lists, we randomly selected the 'noun' form of one of the two lexically-specified sentence variants within a given context type and the 'verb' form of the other sentence variant (e.g., one list may include 9a and 9b and 10a and 10b, but not two minimally different sentences such as 6a and 6b or 7a and 7b), resulting in a list containing 24 noun and 24 verb sentences. We then created a second list in which the grammatical category for each word pair was switched between the sentence variants.

We further manipulated the length of the PP following the target word by assigning the long versions to half of the word pairs in context A and the short versions for the other half. For context B we selected the opposite length setting for the corresponding word pair. This kept the length of the following PP stable between the noun and the verb reading of a given word pair within the same context type, while varying it across the types of sentence context. We then created a mirror-image of these assignments by switching long versions to short and vice versa.

Finally, we merged the two category and sentence variant lists with the length assignments in a 2 x 2 design, resulting in four different lists for reading in which each word pair occurred four times: twice as a noun, twice as a verb; twice in the A context, twice in the B context; twice preceding a short PP, and twice preceding a long PP. Every participant read one list. We pseudo-randomly ordered the paragraph-embedded target sentences in the lists so that at least five other items separated the reading of the same word pair. We also added 24 filler paragraphs, which were also three sentences in length, but contained different, unrelated linguistic material. These fillers were randomly interspersed into the target paragraphs of the four lists, so that each list consisted of 72 paragraphs in total.

2.2 Participants

Participants were 78 adult native speakers of American English (44 male and 34 female). An additional 4 participants were excluded from analysis due to significant disfluency throughout their recordings. All participants were recruited from introductory psychology courses at North

Dakota State University, gave informed consent in written form, and received course credit for their participation.

2.3 Procedure

2.3.1 Recording

Participants were seated approximately 18 inches in front of a laptop computer connected to a Blue Snowball USB microphone. The microphone was placed between them and the laptop keyboard. Recording for each session was done continuously using the software Audacity (Audacity Team, 2017).

Participants were told that they would see paragraphs, one at a time, on the laptop screen. The paragraphs would start in red and then turn black after 10 seconds. Participants were told to read the text silently to themselves while it was red. Then, when it turned black, they were to read it aloud as naturally as possible, using a conversational pace and volume. They were instructed to stop and start over if they made a mistake. Once they had read the paragraph without error, they were instructed to press the space bar to move to the next item.

2.3.2 Preparation and analysis of data

A first automatic segmentation of the data was carried out using the MAUS forced alignment software (Kisler, Reichel, & Schiel, 2017). In a second step, word and vowel segmentation of the target words was corrected manually by research assistants relying on cues in the waveform and the spectrogram, using Praat (Boersma & Weenink, 2016). All research assistants had experience with phonetic segmentation even before this segmentation task and were additionally trained to do the phonetic segmentation of the target words and sentences over the course of several training sessions. A segmentation manual was compiled by the first author, which covered the particular cues for word and vowel segmentation of the different target words' phonetic characteristics and their environments, including instructions for difficult boundaries.

One such example was the location of the left word boundary of the target word *snack* (n/v) when preceded by another sibilant, as in e.g., *the kids' snack / the kids snack*. In cases where the two words were spoken without a pause in between them and there were no clear cues in the spectrogram or the waveform (e.g., a clear dip in intensity between the two segments), the midpoint of the frication of both sibilants was taken as the left word boundary of *snack* (see left panel of Figure 1 for an example segmentation of such a case). Other difficult segmentation decisions involved cases in which the target word ended in a plosive and was followed by another plosive in the onset of the following word (e.g., in the sequence *dig behind*). In some of these instances the first plosive was not released, making it difficult to determine the boundary between the end of the word and the closure phase of the following plosive. When there were no clear cues in the spectrogram or the waveform, we employed the intensity curve as an additional criterion and, if there was a clearly visible low point in the intensity curve, we set the inter-word boundary at this point (see the right panel of Figure 1 for an illustration of such a case). Boundaries determined in this way were further checked by auditory inspection, as were boundaries between plosives for which the intensity curve was not informative (e.g., when intensity remained equally low for the entire duration of the two segments). The segmented data were analysed using a Praat script that extracted word duration, vowel duration and F0 range of the word (in semitones).

Figure 1

Example segmentations (Left panel: target word *snack* (v), right panel: target word *dig* (v), the curve displayed illustrates intensity)

All readings of the target sentences were coded with regard to the prosodic boundary following the target word, employing the ToBI system (Beckman & Hirschberg, 1994). In this system, the break index signifies the degree of disjuncture between each pair of words within an utterance. The relevant ToBI break indices for the present analysis are 4 = intonational phrase (IP) boundary, 3 = intermediate phrase boundary, 1 = word boundary.

In coding, we based our decisions on the instructions given in Beckman & Ayers Elam (1997) and the lecture notes of Veilleux, Shattuck-Hufnagel, & Brugos (2006). As suggested in these resources, the coding of the boundary strength was carried out relying both on the percept of the degree of disjuncture as well as on a tonal analysis of the F0 contour. The coding was performed in a two-step fashion: first it was decided whether the boundary after the target word was a phrase boundary (break index = 3 or 4) or a mere inter-word boundary (break index = 1). In case of a phrase boundary it was then decided whether the phrase boundary concluded an intonational phrase or an intermediate phrase. In the ToBI framework, an intermediate phrase boundary is marked by a phrase-accent, e.g., (L-), while an intonational phrase boundary is marked by an additional boundary tone, e.g., (L-L%). When the phrase accent and the boundary tone are both H or both L tones, which is often the case in our data, the difference between 3 and 4 should be marked by a more pronounced rise or fall of F0 through the presence of the boundary tone. See the following two waveforms and pitch tracks in Figures 2 and 3 for an exemplification of the difference (produced by two male speakers). The two bottom tiers below the spectrogram illustrate the ToBI tone and break index.

Figure 2
Example of an intermediate phrase boundary (break index = 3, see bottom tier) after the target word *kiss* (n) (the speckles indicate F0)

Figure 3
Example of an intonational phrase boundary (break index = 4, see bottom tier) after the target word *kiss* (n) (the speckles indicate F0)

In the example illustrated in Figure 3, the 4 boundary is additionally marked by a pause. Comparing these two examples shows that in case of the 4 boundary F0 drops to a lower level, indicating the occurrence of a low boundary tone. The difference between 3 and 4 boundaries can also be seen by comparing the drop in F0 at the end of the sentence in Figure 2, which marks a 4 boundary, to the moderate decrease in F0 marking the intermediate phrase boundary after the target word *kiss*. In the data, the tonal marking of the difference via the F0 contour was not always as clear as in the examples provided, so we also relied on our perception of the degree of disjuncture between the target word and the following word, as indicated above.

The entire dataset was coded by the first author. Because knowing about the research

hypotheses may impact data coding, the break index coding was also carried out by a research assistant who was naïve to the hypotheses under consideration. The two coders' agreement rate on the type of boundary was 63.1% (33% baseline), yielding a kappa value of 0.45. When conflating 3 and 4 boundaries, i.e., using a binary measure indicating whether a phrase boundary was heard or not, the agreement rate rises to 78.6% (50% baseline) and a kappa value of 0.57. While these agreement rates indicate only moderate agreement, they are almost identical to the agreement rates for prosodic boundary coding obtained in a recent study on inter-rate agreement using the ToBI system (Breen et al. 2012, p. 293). Crucially, both the first author's and the naïve research assistant's coding show the effect of nouns being followed by stronger prosodic boundaries (see Results section below), so we are confident that the effect is not due to the first author's knowledge of the hypotheses influencing the coding. All calculations reported in the following sections rely on the research assistant's coding, but the same pattern of results is obtained when using the first author's coding instead.

2.4 *Collecting contextual predictability ratings*

In addition to the variables of interest, the predictability of the words from context may impact the pronunciation of the target words (Bell, Brenier, Gregory, Girand, & Jurafsky, 2009). Contextual predictability could be a confounding variable if there was a systematic difference in predictability of either the noun or the verb in the test sentences. To control for this factor, we estimated the contextual predictability of both the noun and the verb through a cloze test that we ran via Amazon Mechanical Turk. In this test, the paragraph-embedded target sentences with the extended PPs (see examples 6 and 7 above) were presented with a blank instead of the target word. Participants were asked to fill in the word that first came to their mind when reading the sentence. The sentences were divided up so that each participant saw only a subset of the sentences and none of them saw both the noun and the verb form of the same lexically-specified sentence pair. We collected completions from 40 participants per category-specific sentence.

We measured the predictability of the target words as the ratio of target word responses out of all responses and compared that measure between the two syntactic categories. Inflectional variants of the target word were counted as target word responses. Overall the ratio is very similar between nouns and verbs, with 0.16 for nouns and 0.17 for verbs. This difference is not statistically significant when comparing target word to non-target word responses ($\chi^2(2)=0.55$, $N=3,840$, $p>0.4$). In addition to this global comparison, we also tested for differences in predictability at the level of the individual sentence pairs within the two sentence contexts. Out of 48 pairs, 33 exhibit no significant difference in predictability. Of the 15 pairs for which one of the two words was more predictable, the noun was more predictable in seven sentence pairs and the verb in eight sentence pairs. The cloze test indicates that the sentence pairs in this study were not characterised by a systematic bias in predictability toward one of the grammatical categories. Because contextual predictability may nevertheless affect the pronunciation of the target words, we included the logged ratios of target word responses for a particular sentence as a predictor in the statistical models fitted to word and vowel durations.

3. Results

3.1 *Exclusion of data*

The 78 participants produced 3,744 productions of the target words. Of these, we excluded individual data points when the target word was mispronounced, or when there were severe errors in pronouncing the sentence containing it. These errors include readings in which words in the vicinity of the target word were omitted or replaced, or when the reading of the sentence containing the target word was aborted and then restarted. The data cleaning process resulted in the exclusion of 387 data points, creating a final sample of 3,357 tokens.

3.2 The prosodic boundary following the target word

The first hypothesis we test is whether there is a tendency for nouns to be followed by stronger prosodic boundaries. Figure 4 illustrates the distribution of the different ToBI boundary strengths across the two grammatical categories for the two syntactic contexts and the different lengths of the following PP. Examples of the different conditions for the target word pair *chat* are repeated here to facilitate interpretation of the figure (see Appendix for a complete list of target sentences).

Context A – Short PP

(11a) The kids began a chat in front of the museum. (Noun)

(11b) The kids began to chat in front of the museum. (Verb)

Context A – Long PP

(12a) Their friends started a chat right behind the old statue of President Kennedy. (Noun)

(12b) Their friends started to chat right behind the old statue of President Kennedy. (Verb)

Context B – Short PP

(13a) The parents' chat during the play angers Mr. Robinson. (Noun)

(13b) The parents chat during the play and this angers Mr. Robinson. (Verb)

Context B – Long PP

(14a) The students' chat about the quiz on advanced chemistry makes them feel more confident. (Noun)

(14b) The students chat about the quiz on advanced chemistry and this makes them feel more confident. (Verb)

Figure 4

Prosodic boundaries following the target words in sentence context A (left) and B (right) (x-axis and bar height indicate the number of observations in the individual conditions).

The barplot demonstrates that nouns are followed by prosodic boundaries of greater strength across all conditions. This is shown by the higher ratios of intonational phrase boundaries (=4) and the lower ratios of word boundaries (=1) following the target words in the noun contexts. For example, in context A in the condition with a long PP following the target word, the ratio of IP boundaries (=4) is 23.1% after nouns, but only 16.5% after verbs (see the two leftmost bars in Figure 4). While overall there is a tendency for stronger prosodic boundaries in the A-context

than in the B-context, the differences between nouns and verbs show similar patterns in both, with stronger contrasts between the categories in the B-context (compare the left and right part of Figure 4). We calculated chi-square tests comparing the distribution across the three boundary types between nouns and verbs, which indicate significant differences between the two grammatical categories across all conditions ($p < 0.01$ for all pair-wise tests).⁵ We do not find a difference in boundary strength contingent on the length of the PP following the target words, although stronger boundaries would be expected before longer constituents (Watson, et al., 2006).

To further analyse the differences across the individual conditions, we built a binary logistic regression model predicting the boundary strength after the target word using the *glmer* function of the *lme4* package (Bates, Maechler, Bolker & Walker, 2014) in R (R Core Team, 2014). The dependent variable was whether the target word was followed by a prosodic phrase boundary (3 or 4 boundary in ToBI) or a prosodic word boundary (1 boundary in ToBI). We tested grammatical category, type of syntactic context and length of the prepositional phrase as predictors, as well as their interactions. Initially, the random effects structure included random intercepts for speaker and N/V pair, as well as random slopes for grammatical category by these grouping variables. During model fitting we removed the random slopes for grammatical category as it led to non-convergence of the models. We also removed significant fixed-effect predictors and interactions. PP length was found to not significantly improve the model and was therefore removed. See Table 3 for the fixed-effect model output, and the R code of the selected model after having fitted the model as described.

Table 3 Model output of binary logistic regression model fitted to the presence-absence of a prosodic phrase boundary ($N = 3,357$; $p < 0.1$, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$).

R code: boundary ~ category + context + (category: context) + (1 speaker) + (1 N/V pair)			
	Coefficient	Std. error	<i>t</i>
Intercept	0.7750	0.1640	4.726***
Category=verb	-0.4397	0.1059	-4.151***
Sentence context = B	-0.4888	0.1061	-4.608***
Category = verb by sentence context = B	-0.4504	0.1485	-3.033**

The model output indicates that grammatical category and type of syntactic context, as well as the interaction between the two influence the placement of a prosodic phrase boundary. The model was fitted to the presence of a phrase boundary so that negative values indicate a reduced likelihood of a 3 or 4 boundary. The results show that phrase boundaries are more likely in context A and after verbs (see main effects for context and grammatical category). The interaction term further indicates that the effect of verbs being followed by boundaries of weaker strength is more pronounced in the B context compared to the A context.⁶

The ToBI analysis indicates a difference in terms of prosodic boundary strength between nouns and verbs. However, the perception of a prosodic boundary on the part of the coders may have

⁵ We compared the results of the ToBI coding of the first author to the coding of the naïve coder, the latter of which are reported above. The results from both coders show a significant effect of nouns being followed by boundaries of greater strength, with a slight difference in effect size with the category-difference being more pronounced for the author's coding (Cramer's V (author) = 0.191, Cramer's V (naïve coder) = 0.164).

⁶ As suggested by the editor, we calculated the same model on the data points that were transcribed the same way by both coders. This model yields the same results, with the only exception that the interaction term between grammatical category and type of context is only marginally significant ($p < 0.1$), which is most likely due to the lower statistical power resulting from the smaller sample size, compared to the model based on the entire dataset.

been influenced by their knowledge of the typical occurrence of prosodic boundaries, as the perception of prosodic boundaries is influenced by knowledge of their typical distribution (e.g., Buxó-Lugo & Watson, 2016). We therefore complemented the ToBI analysis by considering an acoustic parameter frequently correlated with prosodic boundaries, namely whether a pause followed the target word. We used 180ms as the threshold for a pauses following (Duez, 1982: 13), who, bases threshold values on an analysis of interview speech and finds 180ms to constitute a lower threshold value for silent pauses. Corresponding to the results for prosodic boundary strength, we find that pauses follow nouns more frequently than verbs (χ -squared = 11.9, $df = 1$, p -value <0.001).

3.3 Durational differences between nouns and verbs

The violin plots in Figure 5 illustrate the density distributions and the mean values of the word and vowel durations by grammatical category and sentence context.

Figure 5

Violin plots of word and vowel durations by grammatical category (the dots and error bars indicate mean and standard error).

Overall, we observe that nouns are pronounced with greater duration than verbs. The difference in mean durations between the two categories is significant in all conditions (word durations: context A: $t = 3.73$, $p < 0.01$; context B: $t = 7.51$, $p < 0.01$, vowel durations: context A: $t = 2.17$, $p < 0.05$; context B: $t = 5.55$, $p < 0.01$). Regarding word durations, the size of the durational difference is 13.03ms in context A and 28.4ms in context B. With regard to vowel durations, the difference is 4.6ms in context A and 12.1ms in the context B. The larger durational difference in context B correlates with a greater difference in prosodic boundary strength between nouns and verbs in that syntactic context (see Section 3.2). Whether the difference in prosodic boundary strength accounts for the durational differences between the categories is investigated in the following section.

3.4 Modeling word duration

3.4.1 Model results

In this section we report the results of two mixed-effects models, fitted to the word and the vowel durations of the target words. These models were built employing the *lmer* function of the *lme4* package (Bates et al., 2014) and the *lmerTest* package (Kuznetsova, Brockhoff, Bojesen & Christensen, 2014) in R. The models include grammatical category and the ToBI

boundary following the target word as fixed effects. The effect of frequency was tested by including a logged frequency ratio of the two homophones of each N/V pair as a fixed-effect predictor. This ratio was calculated via the following equation: $\log_{10}(\text{frequency of word} / \text{frequency of homophone twin})$ (see Lohmann, 2018). Frequency values for that calculation were retrieved from the SUBTLEX-US corpus (Brysbaert et al., 2012; Brysbaert & New, 2009), using lemma frequencies for both categories. The type of sentence context (A vs. B), as well as the length of the PP following the target word (short vs. long) are added as further fixed-effect controls. Besides testing the listed variables as main effects, we also test for interactions of the frequency ratio and grammatical category with type of syntactic context. We employ the contextual speech rate as a further control variable, which we calculated as the log-transformed (base 10) number of syllables per second of the paragraph containing the target word, but not including the target word itself. To control for differences in contextual predictability, we enter the log-transformed (base 10) predictability ratios from the rating study (see Section 2.4 above) into the model as a further fixed effect. We also include the gender of the speaker as a fixed-effect control variable and enter the log-transformed (base 10) F0 range in semitones of each token as a further fixed-effect predictor to control for the possibility of different phonetic implementation of pitch accents (see Section 2.1.2). Including this predictor results in the loss of 32 data points as F0 could not be tracked for these tokens. We started out with a random effects structure of the two models that included random intercepts by word pair, speaker and category-specific word as well as random slopes for category and frequency ratio by word pair and speaker. We removed random effects via backward elimination if they did not contribute significantly to model fit. This process resulted in removing the random intercept for the category-specific word, as well as the random slopes for category from both models. The random slope for frequency ratio by word pair was found to improve the fit of the model predicting word duration but not the model predicting vowel duration and was therefore removed from the latter but not the former model. Both models include random intercepts for speaker and word pair. Regarding fixed effects, we kept fixed-effect control predictors in the model even if they did not contribute significantly to model fit, as long as they did not produce effects that ran counter to our expectations (see e.g., Gelman & Hill, 2007). The only variable removed was the speech rate from the model fitted to vowel durations, as its (small and not statistically significant) effect indicated greater vowel durations in faster speech. The dependent variable in both models was transformed employing Box-Cox power transformation (Box & Cox, 1964), using the *boxcox* function of the MASS package (Venables & Ripley, 2010) in R (R Core Team, 2014). *P*-values were obtained by comparing the complete model to a model without the predictor in question and comparing the fit of the two models via a likelihood ratio test. We calculated variance inflation factors for the two models, all of which were <2.5 , indicating low collinearity.

Table 4 Model output of mixed-effect models fitted to the Box-Cox transformed word and vowel durations ($N = 3,325$; $.p < 0.1$, $*p < 0.05$, $p < 0.01$, $***p < 0.001$).**

R code of the model fitted to word duration (<i>lmer</i> function):		
word_duration ~ speech_rate + category * context + ToBI_boundary + PP_length + gender + frequency_ratio + pitch_range + predictability_ratio + (1 speaker) + (1 + frequency_ratio word) ⁷		
R code of the model fitted to vowel duration (<i>lmer</i> function):		
vowel_duration ~ category * context + ToBI_boundary + PP_length + gender + frequency_ratio + pitch_range + predictability_ratio + (1 speaker) + (1 word)		
Random effects	Word duration	Vowel duration
	Variance / SD	Variance / SD
Speaker (intercept), 78 groups	6.510e-04 / 0.02552	2.155e-03 / 0.04642
Word pair (intercept), 12 groups	1.664e-03 / 0.04080	1.290e-02 / 0.11356
Word pair (frequency ratio)	1.586e-05 / 0.00398	
Residual error	2.318e-03 / 0.04815	6.525e-03 / 0.08078
Intercept and fixed-effect coefficients	Word duration	Vowel duration
	Coefficient / Std. Error / <i>t</i>	Coefficient / Std. Error / <i>t</i>
Intercept	2.204 / 0.0176 / 125.375***	2.319 / 0.0342 / 67.899***
Speech rate	-0.0327 / 0.019 / -1.69 .	
Category = Verb	-0.0022 / 0.0024 / -0.933	-0.0031 / 0.0040 / -0.783
Sentence context = B	-0.0144 / 0.0024 / -5.96***	-0.0057 / -0.0040 / -1.421
Category = verb by sentence context = B	-0.0083 / 0.0036 / -2.445*	-0.0174 / 0.0056 / -3.094**
ToBI boundary = 3	0.0452 / 0.0020 / 22.891***	0.0699 / 0.0032 / 21.129***
ToBI boundary = 4	0.0888 / 0.0025 / 35.326***	0.1084 / 0.0033 / 25.764***
PP length = short	-0.0012 / 0.0017 / -0.736	-0.0049 / 0.0281 / -1.734 .
Gender of speaker (male)	0.0108 / 0.0061 / 1.790 .	0.0389 / 0.0101 / 3.545***
Logged frequency ratio	-0.0053 / 0.0016 / -3.333**	-0.0040 / 0.0018 / -2.222*
Logged F0 range of word	0.0188 / 0.0027 / 7.070***	0.0109 / 0.0045 / 2.439*
Logged predictability ratio	-0.0177 / 0.0112 / -1.585	-0.0546 / 0.0186 / -2.930**

The output of the two models (reported in Table 4) shows that the control variables speech rate, F0 range and contextual predictability produce effects in the predicted directions, as faster speech, lower F0 range and higher predictability are correlated with shorter durations of the target words and vowels. Of these effects, F0 range is statistically significant in both models, while contextual predictability is significant only in the model fitted to vowel durations. Speech rate yielded a marginally significant effect in the model on word durations. Further, we find a significant effect of gender on vowel duration (marginally significant in the model on word duration), with male participants producing tokens with greater vowel duration. This finding stands in contrast with previous research on gender effects on articulation, showing that men tend to speak slightly faster than women (Jacewicz, Fox, O'Neill, & Salmons, 2009). We also find a significant difference between sentence contexts, with significantly shorter word durations in sentence context B as compared to context A.

The strength of the following prosodic boundary as captured by the ToBI index yields a

⁷ The R code provided here does not contain information about the transformation of variable values. As described above, the durations were Box-Cox power-transformed. The scalar predictor variables speech rate, frequency ratio, F0 range of the word, and predictability ratio were log-transformed. Since there were data points for which the variables F0 range and/or predictability ratio assumed a value of 0, which does not allow for log-transformation, a value of 1 was added to all original values of these vectors before log-transformation.

significant effect on the duration of the target words and vowels. In the models reported boundary strength of 1 is taken as the default to which 3 and 4 are compared. The model output indicates that boundaries 3 and 4 are significantly different from 1 and that there is a step-wise increase in pre-boundary lengthening, as can be seen by the positive values for the corresponding coefficient values and the increase in the coefficient value from boundary strength 3 to 4. We built alternative models in which we took boundary strength 3 as the default, which shows that the difference between boundary strength 3 and 4 is also statistically significant in both models.⁸

The predictor ‘category’ does not significantly impact durations as a main effect in either model, but the interaction term of category with type of sentence context does in both models. This interaction indicates that there is no significant effect of category on word and vowel duration in context A, but there is one in context B. Fitting the models to the untransformed durations shows that verbs are predicted to be shorter by 7.1ms in the B context (4.6ms for vowel duration). Taken together, the results indicate that most of the duration difference between nouns and verbs can be accounted for by the type of prosodic boundary following the target words, which reduces the duration difference to practically zero in the A context. It is only in the B context that a small difference remains. We return to this result in Section 4.

3.4.2 *The effect of lexical frequency*

Turning to the influence of lexical frequency, we find that the logged frequency ratio of the noun-verb homophones yields a significant effect both on word and vowel durations. Its negative co-efficient indicates that the more frequent category is pronounced with shorter duration. In order to take a closer look at how category-specific frequency impacts the durations of the individual word pairs, we carried out further analyses.

We fitted a model to word durations with the same fixed effects as reported in Table 4, except for ‘logged frequency ratio.’ That predictor was removed as a fixed effect, as was the corresponding random slope by word pair. Consequently, this model does not contain any information about category-specific frequency. Instead of the frequency-related variables, we inserted a random slope adjustment term of category by word pair. This random slope allows for durational differences between the two homophones within each pair, e.g., that *look*(v) is pronounced with a different duration than *look*(n). Its values indicate differences in duration between the noun and the verb of each N/V pair that cannot be accounted for by other variables in the model. If frequency has a reductive effect on word duration, the values of the random slope adjustment should show that verbs of verb-dominant pairs are shorter in duration than their noun counterparts, while verbs of noun-dominant pairs should be greater in duration than their noun counterparts. For example, because *look*(v) is more frequent than *look*(n), the former is expected to be pronounced with a shorter duration. Conversely, because *book*(v) is less frequent than *book*(n), the former is expected to be pronounced with a greater duration. We extracted the corresponding slope adjustment values and plotted them by the logged noun frequency ratio of the N/V pairs, which indicates differences in frequency between the two words of the pairs (for an explanation of this ratio see Section 2.1.1 and Table 1).

⁸ Taking up the suggestion by an anonymous reviewer, we calculated an additional model that conflates the ToBI boundaries 3 and 4, i.e. employs only a binary distinction of either a prosodic phrase boundary present or absent. This model has a considerably worse model fit than the one reported here, showing that the distinction between the 3 and the 4 boundary is warranted by the data.

Figure 6

Random slope adjustment for grammatical category=verb by word pair plotted against the logged noun ratio of the N/V pairs. (On the x-axis, negative values indicate a higher frequency of the verb, positive values indicate a higher frequency of the noun. On the y-axis negative values indicate a shorter duration of the verb, positive values indicate a shorter duration of the noun.)

Figure 6 demonstrates the expected correlation between the two variables ($r_{\text{pearson}}=0.72$). A model regressing the random slope adjustment on the logged noun ratio of the word pair (see regression line in Figure 6) yields a highly significant effect of the logged noun ratio on the random slope adjustment ($\beta=0.0081$, $t=3.262$, $p<0.01$). These results mean that durational differences between the verb and the noun of each homophone pair are indeed correlated with the difference in frequency between the two words, with the more frequent word being pronounced with shorter duration. To facilitate interpretability of the size of the effect, we calculated the random slope by category also for a model fitted to the untransformed word durations in milliseconds. The following table provides the results of this calculation. It shows random slope adjustments for the category verb for each word pair in milliseconds as an indicator of durational differences, along with the log-transformed frequency difference per word pair (repeated from Table 1 above). In addition, we calculated a relative measure of durational differences by calculating the percentage by which the verb pronunciation is lengthened or shortened relative to the noun pronunciation by the adjustment term. For this measure we divided the random slope adjustments by the mean observed duration of the word pair.

Table 5 *Durational differences and differences in frequency by word pair*

N/V pair	<i>push</i>	<i>look</i>	<i>dig</i>	<i>kick</i>	<i>kiss</i>	<i>pack</i>	<i>cook</i>	<i>chat</i>	<i>tip</i>	<i>nap</i>	<i>snack</i>	<i>book</i>
Logged noun ratio	-1.26	-1.10	-0.74	-0.65	-0.34	-0.25	-0.24	0.12	0.28	0.84	1.06	1.19
Duration adjustment for category = verb (in ms)	-1.37	-13.56	-5.06	-1.87	3.39	-2.73	-3.95	2.71	6.65	1.72	-0.04	14.12
Percentage of duration adjustment for category = verb	-0.46	-5.89	-1.98	-0.70	1.13	-0.87	-1.50	0.83	2.41	0.53	-0.01	5.21

The results show that the durational differences that are due to differences in frequency are rather small, with a durational difference of close to 6% (or 13.6 ms) for the pair *look* constituting the largest effect. The small effect size observed for many pairs is not unexpected, however, as the size of the frequency difference is also fairly small with many pairs. The important point the analyses illustrate is the consistency of the relationship between differences in frequency and differences in duration, which is expressed by a correlation of $r_{\text{pearson}}=0.703$ between the duration adjustment (in ms) and the logged noun ratio. The overall size of the effect is similar to effect sizes observed for spontaneous speech. In the speech corpus dataset analysed in Lohmann (2018), a one-step difference in frequency on the $\log(10)$ -scale resulted in a durational difference of 12.9ms for the monosyllabic word pairs. This is very close to the effects for *look* and *book* in the present data, the two word pairs that are characterized by almost the same frequency difference.

4. Discussion

The present study set out to evaluate the sources of durational differences between nouns and verbs by investigating the acoustic realization of noun-verb homophones in controlled sentence-medial contexts. More specifically, we tested for effects of a category-specific tendency for prosodic boundary placement and an effect of lexical frequency. We find that both of these factors independently influence the duration of noun-verb homophones.

The first hypothesis that this study tested was whether a category-specific difference in strength of the following prosodic boundary could be found. The ToBI coding of prosodic boundary strength shows that nouns are followed by stronger prosodic boundaries than verbs, corroborating the finding by Watson et al. (2006). This finding allowed us to evaluate the hypothesis that the strength of the prosodic boundary following a word accounts for the category-level differences in noun-verb duration. Indeed, the acoustic analysis of the target words robustly shows that verbs are pronounced with shorter word and vowel duration than nouns across all conditions.

Mixed-effects models predicting word and vowel durations make clear the source of these differences, as the models reveal that most of the durational differences between the categories, in fact, are due to the contrast in prosodic boundary strength following the target words. When the ToBI break index of the following prosodic boundary is included as a predictor in the models, the durational differences between the categories are accounted for. The only statistically significant difference between the noun and verb categories that persists is that of a small durational difference in context B. This durational difference could in principle be interpreted as a direct effect of grammatical category on acoustic realization. We consider this unlikely, however. The main reason for this assessment is that we do not find such an effect in context A, which is the syntactic environment that is more carefully controlled between the noun and verb condition and therefore better suited to bring out direct effects of grammatical

category on duration, if they exist. While in context B, which was modelled after the test sentences by Watson et al. (2006), the nouns and verbs occur in similar contexts, these contexts were not syntactically matched (see Section 1.1). The syntactic differences between the two conditions increase the likelihood of prosodic differences between the noun and the verb sentences that may not have been fully captured by the ToBI coding of prosodic boundary strength following the target word. This may have resulted in models that fall short of a complete account of the durational differences between the grammatical categories in that context (see also the discussion of the intonational structure of the test sentences at the end of Section 4.4).

With regard to the hypothesis of a reductive effect of lexical frequency, we found that the higher frequency category is pronounced with shorter duration, affecting both word and vowel duration. This result confirms previous corpus-based studies which had identified frequency as a co-determinant of durational differences between noun-verb homophones (Conwell, 2017; Lohmann, 2018).

In sum, the findings provide evidence for a more complex scenario of what underlies the durational differences between nouns and verbs than had been assumed in previous research, in which differences were mostly attributed to only one factor. The present results indicate that effects of grammatical category on pronunciation are due to the interplay of several variables, which impact acoustic realization via different processes. We turn now to a discussion of how those processes play out to produce differences in the duration of nouns and verbs.

4.1 Nouns and verbs and prosodic boundary placement

This study tested the hypothesis of whether prosodic boundary placement is sensitive to grammatical category. We obtained evidence for such an effect, as across all conditions, prosodic boundaries after nouns were of greater strength than those after verbs. More specifically, when comparing the two sentence contexts that we tested, this effect is more pronounced in the context employed by Watson et al. (2006) (context B, above), in which the target nouns were embedded in a more complex syntactic constituent than the target verbs (see Section 1.1, examples 3a vs. 3b). Still the effect of grammatical category on prosodic boundary placement was also significant in the context in which we controlled for this difference (sentence context A, above).

These findings have implications for accounts of intonational phrasing such as the LRB algorithm by Watson & Gibson (2004). First, the finding of a difference in effect size between the two syntactic environments suggests that the differences in constituent complexity between the noun and the verb condition has had an independent effect on the size of the difference between nouns and verbs. Recall that the length of the embedding clause was not controlled for by Watson et al. (2006) and the complexity of the embedding phrase was the same only when measuring the number of phonological phrases but not when counting the number of words or syllables. While the findings do not indicate whether only one or both of these factors contributed to a stronger boundary strength after nouns, they suggest that in measuring constituent complexity, a more fine-grained measure, such as number of syllables, may be empirically more adequate. Which measure proves to be the best operationalization is an empirical question, whose investigation would be a worthwhile enterprise to refine existing algorithms of prosodic phrasing (e.g., Gee & Grosjean, 1983; Watson & Gibson, 2004).

The second major finding of the present study is that even when carefully controlling constituent structure between the noun and the verb condition, the strength of the prosodic boundary after nouns is still greater than after verbs. As outlined in the introduction, none of the currently existing accounts differentiate between nouns and verbs in predicting the prosodic boundary following these grammatical categories. Implementing such an effect would necessitate an algorithm that treats nouns and verbs differently based on category membership

rather than just based on constituent complexity. A possible implementation within the algorithm of Watson & Gibson (2004, 2005) would be to add a certain additional value indicating the predicted prosodic boundary strength to all phonological phrases ending in nouns (see Section 1.1 on the LRB algorithm). It remains a task for future research to empirically investigate how large this value should be. Addressing this question necessitates the investigation of a variety of different syntactic environments, extending the admittedly very selective testing of few syntactic contexts in the present study.

Regardless of its possible implementation, the category-specific difference in prosodic phrasing raises the question of its cause. A possible explanation may be that the effect is a prosodic generalization effect caused by distributional differences between nouns and verbs with regard to following syntactic constituents. This explanation builds on ideas discussed in Watson et al. (2006). In investigating the effect of the obligatoriness of a following constituent on intonational boundary placement, they suggest that - instead of a binary distinction of obligatory vs. non-obligatory - speakers rely on probabilistic knowledge of the co-occurrence of words with dependent elements, which can e.g., be measured via corpus frequencies (Watson et al. 2006, p. 1052, also Turk 2010: 242-243).

Such probabilistic knowledge may impact prosodic boundary placement, in the sense that a prosodic boundary is more likely after a given word in case the probability of a following dependent constituent is low. Given the general positional differences between nouns and verbs, one may assume that the likelihood of a dependent constituent occurring after verbs is higher than after nouns, while conversely it is more likely that a major syntactic constituent, e.g., a clause or sentence, ends after nouns. We undertook a small-scale corpus analysis to test whether such a distributional difference may explain the pattern in our data. We conducted a lemma search for the twelve target noun-verb pairs in the Contemporary Corpus of American English (COCA) (Davies, 2014) and used punctuation marks in the corpus as an (admittedly coarse) proxy for a major syntactic constituent ending after the word. For each lemma, we calculated the ratio of it being followed by a period or a comma and found a highly significant difference between the noun and verb lemmas, with the average noun being followed by a punctuation mark in 27% of the cases, compared to only 9% for the average verb ($\chi^2=100822.8$, $df=2$, $p<0.001$). We also carried out an analysis on the level of the individual lemma by calculating the correlation coefficient of the ratio of a punctuation mark following the target lemma and the average prosodic boundary strength per lemma, which we arrived at by calculating the mean of the ToBI break scores. This calculation yields an appreciable and statistically significant positive correlation ($r=0.46$, $p<0.05$). These corpus findings indicate that the differences in prosodic boundary strength that we find in our data reflect the occurrence/non-occurrence patterns of the target words with following dependent constituents in natural discourse.

Taking up a suggestion from an anonymous reviewer, we also tested whether the probability of a prepositional phrase following the target words may explain the differences in boundary strength. We investigated this by conducting two analyses. First, we investigated whether there are differences in the general likelihood of nouns and verbs being followed by PPs. We employed the Buckeye corpus (Pitt et al., 2007) to calculate the ratios of nouns and verbs being followed by prepositions. This analysis indicates that there is hardly a difference, with verbs being followed by prepositions in 13.6% and nouns in 11.1% of the cases. We conducted a further investigation in which we investigated the likelihood of a following PP for the individual lemmas. We calculated the ratio of a preposition following the target noun and verb lemmas, employing the COCA corpus. This analysis indicates a slightly lower likelihood of our target verbs being followed by PPs, with ratios of 0.21 for the average verb and 0.27 for the average nouns. This trend would predict that the target nouns would be followed by weaker prosodic boundaries than the target verbs, contra our findings. Calculating the correlation coefficient of the PP ratios with the average ToBI break score for each lemma yields a weak negative correlation ($r=-0.14$). The combined findings of our corpus analyses indicate that it is the general likelihood of being followed by a dependent syntactic element, rather than the

likelihood of a specific phrase type, that determines differences in prosodic boundary strength.

The findings suggest that in prosodic phrasing speakers rely on probabilistic knowledge about the typical environment words occur in. This may be an important factor in explaining the difference in grammatical category that we find in our data: Since the verb lemmas in our dataset are more frequently followed by dependent elements, speakers may show a greater tendency to integrate a following syntactic constituent prosodically. This explanation is compatible with the account by Turk (2010), which argues that prosodic boundary placement is contingent on the relative predictability of linguistic constituents.

An important question is to what extent this tendency is of a language-specific or universal character. The relevant typological dimension is the dominant order of grammatical constituents of a language. In languages typologically different from English, with OV word order and predominantly postposing adpositions, the occurrence bias with regard to following dependent elements could be reversed, with the possible result of weaker prosodic boundaries after nouns than verbs. Investigating this question and the impact of this factor on the acoustic realization of nouns and verbs is an important goal for future research. A challenge in pursuing this goal is that even the universality of the noun-verb distinction as such is contested (see e.g., Haspelmath, 2012).

4.2 Nouns and verbs and frequency-induced reduction

The second hypothesis that we tested was that of a reductive frequency effect, predicting a relatively shorter duration for the word of each noun-verb pair that has the higher usage frequency. The empirical analysis provides evidence for this effect, with the size of the effect depending on the size of the difference in frequency between the noun and the verb. The current finding thus corroborates the results from Lohmann (2018) and Conwell (2017), which were based on analyses of corpus data. It is important to note that while this effect is located at the level of the individual word it does have implications for durational differences at the category level, due to the general difference in type-token distribution between nouns and verbs in discourse (see Section 1). Because verbs have a higher lexical frequency in natural discourse than nouns on average, the net result of this effect will be shorter durations of verbs.

Reductive effects of frequency are well reported, including phonetic effects on the duration of words (see Jurafsky, 2003). How does this effect come about? In previous research on the pronunciation of nouns and verbs, Wang and Xu (2017) equated the frequency effect with a difference in contextual predictability or information load between nouns and verbs, taking up an idea from Coker et al. (1973). However, other studies have shown that the reductive effect of frequency on word duration exerts an influence that is independent of effects of contextual predictability (Gahl, 2008; Lohmann, 2018). Also the present study indicates that frequency yields an independent effect, as contextual predictability was measured separately (see Section 2.4) and, although it impacts duration, the frequency effect persisted even in models including contextual predictability as a control variable.

The precise mechanism of an independent frequency effect has been debated in the literature, with particular focus on the stage of language production in which it occurs. Two important positions are that it is regarded either as a purely post-lexical effect of articulatory automatization, or as an effect that speaks to differences in representation between low-frequency and high-frequency words (see discussions in Gahl, 2008; Jurafsky, 2003). Studies that show that the frequency effect persists in homophones, like the present one or Gahl (2008), provide an argument against an explanation in terms of articulatory automatization, which would predict the shortening of words to be the result of form frequency. Homophones share the same form on the segmental level and should therefore be subject to the same effects of articulatory practice, which evidently is not the case. Consequently, accounts have been put forth that argue the phonetic shortness of frequent words to be the result of differences between

the lexical entries of high-frequency and low-frequency words. One explanation rests on the assumption that low-frequency words have a lower resting activation than high-frequency words and are therefore more slowly retrieved. The slower retrieval of low-frequency words is reflected in greater durations in order to adjust for asynchronies between planning and implementation (see Bell et al., 2009). A different approach is to locate differences in duration as reflecting differences in the storage of phonetic detail. A class of models pursuing this approach are exemplar-based models, which allow for the representation of word-specific phonetic detail. Within an exemplar-based model, each word is represented by many exemplars forming an exemplar ‘cloud’ that emerges via the processing of individual tokens. The exemplar ‘clouds’ of high and low-frequency words would be expected to differ, with a greater number of shorter tokens in the high-frequency group and a greater number of longer tokens in the low-frequency group (see Pierrehumbert (2001, 2002) on how frequency effects are modelled in exemplar-based models). With regard to general, phonetic differences between nouns and verbs, such accounts would predict that most verbs are represented by acoustically shorter exemplars than most nouns, because on average the usage frequency of verbs is higher than that of nouns.

Regardless of the exact implementation, the fact that a frequency effect supports the differentiation of the acoustic forms of noun-verb homophone pairs suggests that the noun and the verb occupy entries in the lexicon with different specifications. This has implications for models of the mental lexicon in linguistics, as well as models of speech production. The finding of a frequency effect differentiating homophones favours models that assume separate representations for the noun and the verb of noun-verb homophone pairs (e.g., *kick(n)/kick(v)*), but not models that assume a shared entry (see the discussion in Lohmann, 2018). Furthermore, following Gahl (2008), the frequency effect on homophones runs counter to assumptions of speech production models in which homophones share the same phonological form entry, most importantly the model by Levelt and colleagues (Levelt, 1989; Levelt, Roelofs, & Meyer, 1999). Within this model, frequency is ‘inherited’ between homophones at the level of the shared phonological form; the two words should therefore be pronounced with equal duration. Models that can accommodate the frequency effect are those that do not assume frequency inheritance and argue for completely separate representations of homophones (Caramazza, Costa, Miozzo, & Bi, 2001; Miozzo & Caramazza, 2005).

4.3 Can the effects be explained in a unified account?

The present reading study was designed to isolate the impact of lexical frequency and prosodic factors on the durational characteristics of noun-verb homophones. In doing so it pursued an approach that disentangles variables that can be assumed to be strongly correlated in discourse. The opposite of such an approach would be to develop a unifying account that tries to integrate both groups of factors. One such account is the *Smooth Signal Redundancy Hypothesis* (SSRH) by Aylett & Turk (2004). This account points to the strong relationship between language redundancy, that is, the predictability of a given word or constituent, and acoustic redundancy, that is, the salience of the corresponding acoustic signal. It states that the more predictable or redundant a linguistic element is, the less salient its acoustic realization will be. This lower salience may, for example, be signalled via a reduced duration. In that account, the prosodic structure of the sentence acts as an intermediary. Hence, redundancy does not directly impact phonetic implementation, but does so only via variations of the prosodic structure of the sentence (see also Aylett & Turk 2004: Figure 2). While the original formulation of the SSRH dealt only with differences in prosodic prominence, Turk (2010) argues that also the placement of prosodic boundaries is influenced by differences in redundancy. Given the obvious relevance of this account for the present results, we would like to discuss to what extent our findings are compatible with the SSRH.

We consider our results to be largely in line with the SSRH. The finding that the target verbs are followed by prosodic boundaries of lesser strength than the nouns confirms assumptions of the

SSRH. Our corpus analysis shows that dependent linguistic material is more likely to follow the verbs, resulting in a difference in predictability and therefore redundancy of the following dependent constituents between categories. Ergo, a weaker prosodic boundary would be expected after verbs within the SSRH framework. By showing that the difference in prosodic boundary strength impacts the duration of the target words, our study shows that linguistic redundancy impacts prosodic structure, which then influences acoustic realization, as predicted by the SSRH (see in particular Turk 2010: 240-242).

However, the effect of lexical frequency may be more difficult to explain within the framework of the SSRH. In the literature on the SSRH the usage frequency of a word is usually interpreted as contributing to its general predictability (see the discussion in Section 4.2, see also Turk 2010: 252), which renders it a measure of language redundancy. The finding of a frequency-induced reduction effect is therefore in line with the general assumption of a link between linguistic and acoustic redundancy made by the SSRH. However, it is important to note that we find the effect of frequency-induced reduction even after controlling for prosodic effects. The correlation illustrated in Figure 6 shows the reductive effect of frequency on word duration after controlling for the strength of the prosodic boundary following the target words. This finding is not straightforwardly explained by the SSRH, as the SSRH assumes that redundancy effects are mediated by prosody. It is important to note, however, that, corresponding to our results, Aylett & Turk (2004, p. 53) and also Baker & Bradlow (2009) found small independent effects of redundancy when prosodic structure was controlled for. One possibility is that this pattern of results comes about through an imperfect operationalization of prosodic factors (see the discussion in Aylett & Turk (2004, 52)). In the present case, it is possible that the ToBI measurement of boundary strength is not fine-grained enough to control for the impact of differences in prosodic boundary strength in their entirety, so that residual domain-final lengthening effects were picked up by the frequency measure we employed. However, given that ours is not the only study to report independent effects of redundancy, it is also possible that these truly are independent. This possibility raises the question of how this independence is best implemented in the SSRH, a question that is beyond the scope of the present paper, but may be an interesting topic for future research.

4.4 Caveats of the present study and avenues for future research

The present study provided evidence that durational differences between noun-verb homophones are co-determined by more pronounced prosodic boundaries after nouns than verbs, and a reductive effect of token (lemma) frequency. We must keep in mind, however, that the present study did not control for all factors possibly co-varying between the two grammatical categories. Additional investigation may therefore shed more light on the source of difference in acoustic realization between the two grammatical categories. The target sentences were designed to hold the impact of syntactic constituent structure on prosodic boundary placement after the target words constant. In addition, we aimed to control the overall prosodic characteristics of the sentences between the noun and the verb condition, at least in the environment designed by ourselves (context A). However, we cannot be entirely certain that these are in fact identical, as factors such as the overall intonational characteristics of the sentences, measurable as, for example, prosodic breaks in other positions of the sentence, or the particular sequence and distribution of pitch accents across the sentence may differ in subtle but significant ways between the noun and the verb condition. We invite future research to further investigate the impact of such variables in testing our findings.

We would like to single out the pitch accent structure of the target sentences in particular. Our sample analysis reported in Section 2.1.2 does not indicate systematic differences with regard to the occurrence and type of pitch accent on the target words between the noun and the verb condition. However, it remains possible that a more thorough analysis based on a larger sample that also considers the types of pitch accents in the vicinity of the target words may bring to light differences between the two conditions. Niebuhr & Pfitzinger (2010) show that different

types of pitch accent can differ in their particular durational and intensity characteristics, which affect not only the accented, but also neighbouring syllables. This means that beyond the pitch accent on the target word, one would ideally examine further accents in its vicinity. Such an analysis would not just be relevant for the analysis of the durational differences we address in the current paper. It would also contribute to the question of a prominence difference between nouns and verbs that is discussed in intonational phonology (see Ladd, 2008, pp. 246–247), because previous empirical studies on that topic, such as Jensen (2006), do not take into account the entire pitch accent structure of the sentence.

5. Conclusion

This paper provides empirical evidence that differences in duration between nouns and verbs have multiple sources. The findings thus indicate that beyond a difference in sentence positions and concomitant prosodic contrasts, further factors underlie the pronunciation differences between the categories. The two main processes identified are a category-specific propensity for more pronounced prosodic boundaries following nouns vs. verbs even in matched syntactic contexts and a reductive effect of lexical frequency of the category-specific word.

While painting a more complex picture of the factors underlying acoustic effects of grammatical category, the results agree with the conclusion drawn in previous research that grammatical category does not directly influence speech timing (Cooper & Paccia-Cooper, 1980; Sorensen et al., 1978). Acoustic differences come about indirectly via the interplay between grammatical category and the processes of prosodic phrasing and frequency-induced reduction. Due to the distributional characteristics of nouns and verbs in naturally occurring speech, these processes will in many cases team up and produce the durational differences that have been widely observed.

Acknowledgements

We thank the audience at the workshop of the research group ‘spoken morphology’ at HHU Düsseldorf in August 2017 and the audience at the Alberta Phonetics Laboratory in Edmonton in September 2017 for valuable feedback. We express our gratitude to our research assistants, Viktoria Schneider, Elena Girardi, Mandy Rilke, Samantha Hamernick, Rachel Helgeson, Amanda Auen, and Allyson Kuznia who helped us with the recording, segmentation, and coding of the data. This work was supported by the *Deutsche Forschungsgemeinschaft* under grant LO-2135/1-1 (to AL) and the National Institute of Child Health and Human Development (NICHD) under grant 1R15HD077519-01 (to EC). The contents of this article are the sole responsibility of the authors and do not necessarily reflect the views of NICHD.

Competing interests statement: The authors have no competing interests to declare.

Appendix

Paragraphs and sentences of sentence context A:

The material in parentheses represents the sentence version with long/extended PP.

chat (N/V)

(1)

On Sunday the Pattersons spent a relaxed afternoon in the city center. John and Emily were having ice cream.

Noun sentence: The kids began a *chat* in front of the museum (of Natural History).

Verb sentence: The kids began to *chat* in front of the museum (of Natural History).

(2)

Helen and Gabe were enjoying a great time in the historic part of town, where they also met Michael and Sophia. Helen and Gabe bought postcards for their parents.

Noun sentence: Their friends started a *chat* right behind the old statue (of president Kennedy).

Verb sentence: Their friends started to *chat* right behind the old statue (of president Kennedy).

cook (N/V)

(1)

Everyone who works at the law firm does something interesting after hours. Melody grows flowers and hopes to start a florist shop.

Noun sentence: Benedict is a *cook* after his daytime job (at the office).

Verb sentence: Benedict wants to *cook* after his daytime job (at the office).

(2)

With lots of visiting family, everyone will need to help out. There's just too much for two people to do.

Noun sentence: Jill and Mike need a *cook* for the first seven days (of September).

Verb sentence: Jill and Mike need to *cook* for the first seven days (of September).

dig (N/V)

(1)

People of all ages came to the beach to play. The teenagers played volleyball, while the toddlers napped.

Noun sentence: The little children started a *dig* before dinner (with their parents).

Verb sentence: The little children started to *dig* before dinner (with their parents).

(2)

A famous Spanish pirate was rumored to have buried his treasure near the ocean. After discovering an old map with a big X, Ricki thought she knew where to look.

Noun sentence: The treasure hunter began the *dig* behind the rock (on the beachfront).

Verb sentence: The treasure hunter began to *dig* behind the rock (on the beachfront).

book (N/V)

(1)

If you don't know the story before you see the opera, you'll probably be confused. If you do know it, with the right seats, it's an incredible show.

Noun sentence: I kindly suggest you read this *book* for the performance (on Sunday night).

Verb sentence: This is the ticket you want to *book* for the performance (on Sunday night).

(2)

When you're preparing for your trip, you'll want to make sure you know everything about Taiwan and how to get there. I may have some useful advice for you.

Noun sentence: I kindly suggest you scan this *book* with the new iPad (on the table)

Verb sentence: This is the cheap flight you want to *book* with the new iPad (on the table).

kick (N/V)

(1)

The Italian team had fallen behind because their forward wasn't playing very well. The coach was getting angry with him.

Noun sentence: He gave him a *kick* after the whistle (of the referee).

Verb sentence: He asked him to *kick* after the whistle (of the referee).

(2)

Too bad Frank and Tom had forgotten the key to their old home. Frank tried to push the door open, but it didn't move.

Noun sentence: Tom gave it a *kick* with his new boots (of hardened leather).

Verb sentence: Tom asked him to *kick* with his new boots (of hardened leather).

kiss (N/V)

(1)

In our family, everyone has a different favorite Christmas tradition. Tim and Andrea love making new ornaments to decorate the tree.

Noun sentence: Jill gives Martin a *kiss* under the mistletoe (above the door).

Verb sentence: Jill asks Martin to *kiss* under the mistletoe (above the door).

(2)

Natalie wants to go into the city, but she doesn't really want to go with her friends. The last time she went with them, she was really embarrassed. *Noun*

sentence: Helen gave Bart a *kiss* on the underground train (to Manhattan).

Verb sentence: Helen asked Bart to *kiss* on the underground train (to Manhattan).

look (N/V)

(1)

The children all thought that the house at the end of the street was haunted. They drew straws to see who would go peek inside.

Noun sentence: Florence had a *look* through a window (of the old house).

Verb sentence: Florence had to *look* through a window (of the old house).

(2)

After they redecorated the conference room, no one in the office could stop talking about it. The staff all thought the new color scheme was ridiculous.

Noun sentence: Mary had a *look* during her break (of ten minutes).

Verb sentence: Mary had to *look* during her break (of ten minutes).

nap (N/V)

(1)

After the PE class, both students and teachers were really tired. They were looking for a place to rest. The teachers went to the cafeteria.

Noun sentence: The students had a *nap* in the classroom (next to the gym).

Verb sentence: The students had to *nap* in the classroom (next to the gym).

(2)

Brian, the cook at the daycare, often prepares huge portions for the kids. Because of that they get really tired – I saw that on Monday.

Noun sentence: The children had a *nap* after their lunch (of hamburgers).

Verb sentence: The children had to *nap* after their lunch (of hamburgers).

pack (N/V)

(1)

Robin went hiking on the Annapurna trail in Nepal. On the second day, the group took a rest in front of a temple, before continuing their ascent.

Noun sentence: The young sherpa shouldered the *pack* during the break (before the climb).

Verb sentence: The young sherpa began to *pack* during the break (before the climb).

(2)

In preparation for the Obamas' trip to Europe some of their staff had to work overtime. Many pieces of luggage had to be prepared, including a huge suitcase.

Noun sentence: Their old nanny carried the *pack* for some three hours (of her free time).

Verb sentence: Their old nanny needed to *pack* for some three hours (of her free time).

push (N/V)

(1)

At a board meeting, the executives of the company discuss their strategy for the next year. They decide to invest in more advertising in the Southwest.

Noun sentence: This is the market that needs a *push* over the next quarter (of the fiscal year).

Verb sentence: This is the market they need to *push* over the next quarter (of the fiscal year).

(2)

Chris and Robert are planning the expansion of their company. They decide to promote their newest app.

Noun sentence: This is the product that needs a *push* for their market entry (in South Africa).

Verb sentence: This is the product they need to *push* for their market entry (in South Africa).

snack (N/V)

(1)

Robert, the new representative of the labor union, asks what the factory employees want to make their jobs better. They tell Robert that they need tables in the break rooms.

Noun sentence: The workers like a *snack* during their break (of twenty minutes).

Verb sentence: The workers like to *snack* during their break (of twenty minutes).

(2)

After half an hour the girl scouts took a break to eat. It had rained a lot, so they couldn't sit in their favorite spot under the tree.

Noun sentence: The girl scouts had a *snack* behind the barn (next to the orchard).

Verb sentence: The girl scouts had to *snack* behind the barn (next to the orchard).

tip (N/V)

(1)

Helen and Florence were on a trip in France and didn't understand local customs. They wondered why their server was so angry with them after they left the bistro.

Noun sentence: The waiter expected a *tip* after the meal (in a restaurant).

Verb sentence: The waiter wanted them to *tip* after the meal (in a restaurant).

(2)

Pablo wonders why the staff gets anxious right around the start of dinner service. Rodrigo, the restaurant manager, has an idea.

Noun sentence: Many servers received no *tip* before the break (at seven o'clock).

Verb sentence: Maybe customers need to *tip* before the break (at seven o'clock).

Paragraphs and sentences of sentence context B:

The material in parentheses represents the sentence version with long/extended PP.

chat (N/V)

(1)

Ben and Susan wonder why their teacher always gets aggravated at the theater. They realize it's because of the chaperones' bad behavior.

Noun sentence: The parents' *chat* during the play (on US history) angers Mr. Robinson.

Verb sentence: The parents *chat* during the play (on US history) and this angers Mr. Robinson.

(2)

Ms. Butler, the science teacher, comments to her colleague that her students are very talkative before exams. She suggests that there is a reason for this.

Noun sentence: The students' *chat* about the quiz (on advanced Chemistry) makes them feel more confident.

Verb sentence: The students *chat* about the quiz (on advanced Chemistry) and this makes them feel more confident.

cook (N/V)

(1)

When the children visit their relatives, everything is different. They never know what to expect.

Noun sentence: Their grandparents' *cook* with the bright clothing (from India) entertains Louis and Robin.

Verb sentence: Their grandparents *cook* with special spices (from India) and this delights Louis and Robin.

(2)

The Hendersons were known to be wealthy and flamboyant. They hosted a large party following the annual travel agents' meeting.

Noun sentence: The Hendersons' *cook* for the reception (at the conference) entertains the invited guests.

Verb sentence: The Hendersons *cook* for the reception (at the conference) and this delights the invited guests.

dig (N/V)

(1)

Maria and Pedro had their property landscaped by a garden designer. One day, their neighbor's dogs come through a hole in the fence.

Noun sentence: The dogs' *dig* behind the shed (in the yard) upset Maria and Pedro.

Verb sentence: The dogs *dig* behind the shed (in the yard) and this upsets Maria and Pedro.

(2)

Natalie and Carson are pretending to be archaeologists. They put on pith helmets and took their shovels across the street.

Noun sentence: The kids' *dig* at the playground (in the park) entertains the parents.

Verb sentence: The kids *dig* at the playground (in the park) and this entertains the parents.

book (N/V)

(1)

Some of the hotel guests are confused and don't know how to request rooms. Others leave personal items lying around.

Noun sentence: The visitor's *book* at the front desk (in the lobby) irritates the service staff.

Verb sentence: The visitors *book* at the front desk (in the lobby) and this irritates the service staff.

(2)

People who aren't used to traveling may be unfamiliar with airport security procedures. You can't leave your belongings lying around and you can reserve a ticket only at designated help desks.

Noun sentence: The passenger's *book* at the counter (next to Gate 4) alarms the security guards.

Verb sentence: The passengers *book* at the counter (next to Gate 4) and this alarms the security guards.

kick (N/V)

(1)

Jenny has been working for weeks to improve her offense skills. She is very pleased when she and her teammates are asked to demonstrate that.

Noun sentence: The player's *kick* in front of the cheering group (of spectators) satisfies the coach.

Verb sentence: The players *kick* in front of the cheering group (of spectators) and this satisfies the coach.

(2)

At the beginning of the boxing match the rules are explained. The opponents are allowed to fight each other with their hands only.

Noun sentence: The boxer's *kick* in front of the referee (of the event) upsets the spectators.

Verb sentence: The boxers *kick* in front of the referee (of the event) and this upsets the

spectators.

kiss (N/V)

(1)

During halftime of the game, the most popular students are announced as part of the Homecoming court. Not everyone is as happy about this as they are.

Noun sentence: The couple's *kiss* in front of the crowd (of football spectators) irritates some of the older people.

Verb sentence: The couples *kiss* in front of the crowd (of football spectators) and this irritates some of the older people.

(2)

The gossip surrounding the famous couple has been building for weeks. Everyone who interacts with them is getting really tired of it.

Noun sentence: The actors' *kiss* on the movie set (for the new production) annoys the director.

Verb sentence: The actors *kiss* on the movie set (for the new production) and this annoys the director.

look (N/V)

(1)

At the premiere of the new play Steve manages to sneak behind the stage. From his spot in the corner he witnesses an argument between the director and some of the actors.

Noun sentence: The actors' *look* through the curtains (of the theater) irritates the director.

Verb sentence: The actors *look* through the curtains (of the theater) and this irritates the director.

(2)

Sharyl started a new job last week, but she's still getting used to her gossipy new colleagues. Today they have their first department-wide meeting.

Noun sentence: Her colleague's *look* across the room (on the second floor) irritates Sharyl.

Verb sentence: Her colleagues *look* across the room (on the second floor) and this irritates Sharyl.

nap (N/V)

(1)

Mike and his team are very busy finishing up the report for the end of the quarter. They see that some of their co-workers in accounting do not seem to take their work seriously.

Noun sentence: Their colleagues' *nap* in the cubicle (next to the busy hallway) upsets the hard-working employees.

Verb sentence: Their colleagues *nap* in the cubicle (next to the busy hallway) and this upsets the hard-working employees.

(2)

Dr. Butler and Dr. Gonzales have moved their practice out of the city. Now, some of the older patients are very sleepy when they arrive at the cardiologists' new office.

Noun sentence: The patients' *nap* in the waiting room (with the new furniture) irritates the doctors.

Verb sentence: The patients *nap* in the waiting room (with the new furniture) and this irritates the doctors.

pack (N/V)

(1)

Peter and JJ were playing by the school when some dark clouds rolled in. Their mother had told them to keep their things inside in case of rain, but they didn't listen.

Noun sentence: The young boys' *pack* under the tree (near the playground) got wet in the rain.

Verb sentence: The young boys *pack* under the tree (near the playground) and get wet in the rain.

(2)

After they found their cabin, Barb and Todd began getting ready for the next day. They wanted

to get an early start, so Todd got everything organized.

Noun sentence: The hikers' *pack* for the long hike (in the mountains) was prepared the night before.

Verb sentence: The hikers *pack* for the long hike (in the mountains) and prepare the night before.

push (N/V)

(1)

As globalization progresses, some countries are becoming nervous. This can be witnessed in China.

Noun sentence: The foreign company's *push* into the Chinese market (of mobile computing) alarms the communist party.

Verb sentence: The foreign companies *push* into the Chinese market (of mobile computing) and this alarms the communist party.

(2)

Corporations aren't always concerned with what's best for the Earth. When oil prices are high, they stop at nothing to extract more and more.

Noun sentence: The oil corporations' *push* for extensive investment (in the fracking sector) worries environmentalist groups.

Verb sentence: The oil corporations *push* for extensive investment (in the fracking sector) and this worries environmentalist groups.

snack (N/V)

(1)

A high school teacher is preparing for a long meeting with mothers of her students. She wants to be sure that they don't get too hungry.

Noun sentence: The moms' *snack* at the meeting (on the new school books) is carrot cake.

Verb sentence: The moms *snack* at the meeting (on the new school books) and eat carrot cake.

(2)

Callie and Hunter were watching TV and got hungry. Even though they knew they were going out to eat in an hour, they made themselves peanut butter sandwiches.

Noun sentence: The kids' *snack* before dinner (with their grandparents) makes their parents angry.

Verb sentence: The kids *snack* before dinner (with their grandparents) and this makes their parents angry.

tip (N/V)

(1)

The window-washers are having a hard day because they have made hardly any money. Then an expensive car pulls up and their day gets a lot better.

Noun sentence: The Porsche driver's *tip* at the car wash (on Truman Avenue) pleases the service staff.

Verb sentence: The Porsche drivers *tip* at the car wash (on Truman Avenue) and this pleases the service staff.

(2)

The catering staff is very tired because the holiday season is long and busy. Tonight, they get a great surprise.

Noun sentence: The wealthy couple's *tip* at the party (after the reception) delights the employees.

Verb sentence: The wealthy couples *tip* at the party (after the reception) and this delights the employees.

References

- Aylett, M., & Turk, A. E. (2004). The Smooth Signal Redundancy Hypothesis: A Functional Explanation for Relationships between Redundancy, Prosodic Prominence, and Duration in Spontaneous Speech. *Language and Speech*, 47(1), 31–56.
- Bates, Douglas, Martin Maechler, Ben Bolker & Steven Walker. 2014. *lme4: Linear mixed-effects models using Eigen and S4*. <http://cran.r-project.org/package=lme4>.
- Beckman, M. E., & Ayers Elam, G. (1997). *Guidelines for ToBI labeling: Unpublished manuscript* (Version 3.0): Ohio State University.
- Beckman, M. E., & Hirschberg, J. (1994). *The ToBI annotation conventions*. Columbus, Ohio: Ohio State University.
- Bell, A., Brenier, J. M., Gregory, M., Girand, C., & Jurafsky, D. (2009). Predictability effects on durations of content and function words in conversational English. *Journal of Memory and Language*, 60(1), 92–111. <https://doi.org/10.1016/j.jml.2008.06.003>
- Boersma, P., & Weenink, D. (2016). Praat: Doing phonetics by computer. Retrieved from <http://www.praat.org>
- Box, G., & Cox, D. (1964). An Analysis of Transformations. *Journal of the Royal Statistical Society. Series B (Methodological)*, 26(2), 211–252.
- Brysbaert, M., & New, B. (2009). Moving beyond Kucera and Francis: a critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods*, 41(4), 977–990. <https://doi.org/10.3758/BRM.41.4.977>
- Brysbaert, M., New, B., & Keuleers, E. (2012). Adding Part of Speech information to the SUBTLEX-US word frequencies. *Behavior Research Methods*, 44(4), 991–997.
- Buxó-Lugo, A., & Watson, D. G. (2016). Evidence for the influence of syntax on prosodic parsing. *Journal of Memory and Language*, 90, 1–13. <https://doi.org/10.1016/j.jml.2016.03.001>
- Caramazza, A., Costa, A., Miozzo, M., & Bi, Y. (2001). The specific-word frequency effect: Implications for the representation of homophones in speech production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(6), 1430–1450. <https://doi.org/10.1037//0278-7393.27.6.1430>
- Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. *Studies in language*. New York: Harper & Row.
- Coker, C. H., Umeda, N., & Browman, C. P. (1973). Automatic Synthesis from Ordinary English Text. *IEEE Transactions on Audio and Electroacoustics*, AU-21(3), 293–298.
- Conwell, E. (2017). Prosodic disambiguation of noun/verb homophones in child-directed speech. *Journal of Child Language*, 44(3), 734–751. <https://doi.org/10.1017/S030500091600009X>
- Conwell, E., & Morgan, J. L. (2012). Is It a Noun or Is It a Verb? Resolving the Ambicategoricity Problem. *Language Learning and Development*, 8(2), 87–112. <https://doi.org/10.1080/15475441.2011.580236>
- Cooper, W. E., & Paccia-Cooper, J. M. (1980). *Syntax and Speech*. Cambridge, Massachusetts: Harvard University Press.
- Davies, M. (2014). The Corpus of Contemporary American English: 450 million words, 1990–2012. Retrieved from Corpus also available online at <http://corpus.byu.edu/coca/>
- Dilts, P. (2013). Modelling phonetic reduction in a corpus of spoken English using Random Forests and Mixed-Effects Regression (PhD Thesis). University of Alberta, Edmonton, Alberta.

- D'Imperio, M., Giorka, E., Frota, S., Prieto, P., & Vigário, M. (2005). Intonational phrasing in Romance: The role of syntactic and prosodic structure. In M. J. Freitas, S. Frota, & M. Vigário (Eds.), *Phonology and Phonetics [PP]: Vol. 9. Prosodies: With Special Reference to Iberian Languages* (pp. 59–98). Berlin, Boston: De Gruyter Mouton.
- Duez, D. (1982). Silent and Non-Silent Pauses in Three Speech Styles. *Language and Speech*, 25(1), 11–28. <https://doi.org/10.1177/002383098202500102>
- Ferreira, F. (2007). Prosody and performance in language production. *Language and Cognitive Processes*, 22(8), 1151–1177. <https://doi.org/10.1080/01690960701461293>
- Frota, S., D'Imperio, M., Elordieta, G., Prieto, P., & Vigário, M. (2010). The phonetics and phonology of intonational phrasing in Romance. In P. Prieto i Vives, J. Mascaró, & M. J. Solé (Eds.), *Amsterdam studies in the theory and history of linguistic science. Series IV, Current issues in linguistic theory: v. 282. Segmental and prosodic issues in romance phonology* (Vol. 282, pp. 131–153). Amsterdam, Philadelphia: John Benjamins Pub. <https://doi.org/10.1075/cilt.282.10fro>
- Gahl, S. (2008). Time and Thyme are not Homophones: The Effect of Lemma Frequency on Word Durations in Spontaneous Speech. *Language*, 84(3), 474–496.
- Gee, J. P., & Grosjean, F. (1983). Performance Structures: A Psycholinguistic and Linguistic Appraisal. *Cognitive Psychology*, 15, 411–458.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel, hierarchical models* (1. publ). *Analytical methods for social research*. Cambridge [u.a.]: Cambridge Univ. Press.
- Haspelmath, M. (2012). How to compare major word-classes across the world's languages. In T. Graf, D. Paperno, A. Szabolcsi, & J. Tellings (Eds.), *Theories of everything: in honor of Edward Keenan: UCLA Working Papers in Linguistics 17* (109–130). Los Angeles: UCLA.
- Jacewicz, E., Fox, R. A., O'Neill, C., & Salmons, J. (2009). Articulation rate across dialect, age, and gender. *Language Variation and Change*, 21(2), 233–256. <https://doi.org/10.1017/S0954394509990093>
- Jensen, C. (2006). Are Verbs Less Prominent? *Lund University, Centre for Languages & Literature, Department of Linguistics & Phonetics Working Papers*, 52, 73–76.
- Jurafsky, D. (2003). Probabilistic Modeling in Psycholinguistics: Linguistic Comprehension and Production. In J. Hay, R. Bod, & S. Jannedy (Eds.), *Probabilistic linguistics* (pp. 39–95). Cambridge, MA: MIT Press.
- Kisler, T., Reichel, U., & Schiel, F. (2017). Multilingual processing of speech via web services. *Computer Speech and Language*, 45, 326–347.
- Kuznetsova, Alexandra, Per Bruun Brockhoff & Rune Haubo Bojesen Christensen. 2014. *lmerTest*. <http://cran.r-project.org/web/packages/lmerTest/index.html>.
- Ladd, D. R. (2008). *Intonational phonology* (2nd ed.). *Cambridge studies in linguistics: Vol. 119*. Cambridge, UK, New York: Cambridge University Press. Retrieved from <http://site.ebrary.com/lib/alltitles/docDetail.action?docID=10289418>
- Levelt, W. J.M. (1989). *Speaking: From intention to articulation. ACL-MIT Press series in natural-language processing*. Cambridge, Mass: MIT Press.
- Levelt, Willem J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, 22(01), 1–75. <https://doi.org/10.1017/S0140525X99001776>
- Lightfoot, M. J. (1970). Accent and Time in Descriptive Prosody. *WORD*, 26(1), 47–64. <https://doi.org/10.1080/00437956.1970.11435580>

- Lohmann, A. (2018). cut (N) and cut (V) are not homophones - Lemma frequency affects the duration of noun-verb conversion pairs. *Journal of Linguistics*, 54(4), 753–777. <https://doi.org/10.1017/S0022226717000378>
- Miozzo, M., & Caramazza, A. (2005). The representation of homophones: evidence from the distractor-frequency effect. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 31(6), 1360–1371. <https://doi.org/10.1037/0278-7393.31.6.1360>
- Niebuhr, O., & Michaud, Alexis. (2015). Speech Data Acquisition: The Underestimated Challenge. *Kieler Arbeiten in Linguistik und Phonetik (KALIPHO), Institut für Skandinavistik, Frisistik und Allgemeine Sprachwissenschaft (ISFAS)*, 1–42.
- Niebuhr, O., & Pfitzinger, H. (2010). On pitch-accent identification - The role of syllable duration and intensity. *Proceedings of the 5th International Conference on Speech Prosody*, 1–4.
- Pierrehumbert, J. (2001). Exemplar dynamics: Word frequency, lenition and contrast. In J. Bybee & P. J. Hopper (Eds.), *Typological studies in language: v. 45. Frequency and the Emergence of Linguistic Structure* (pp. 137–157). Amsterdam/Philadelphia: John Benjamins Publishing Company.
- Pierrehumbert, J. (2002). Word-specific phonetics. *Laboratory Phonology. Laboratory Phonology*, 7, 101–139.
- Pitt, M. A., Dilley, L., Johnson, K., Kiesling, S., Raymond, W., Hume, E., & Fosler-Lussier, E. (2007). Buckeye Corpus of Conversational Speech. Columbus, Ohio: Department of Psychology, Ohio State University. Retrieved from www.buckeyecorpus.osu.edu
- R Core Team. (2014). R: A Language and Environment for Statistical Computing}. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <http://www.R-project.org/>
- Selkirk, E. O. (1984). *Phonology and syntax: The relation between sound and structure. Current studies in linguistics series: Vol. 10*. Cambridge, Mass: MIT-Press.
- Selkirk, E. O. (2011). The Syntax-Phonology Interface. In J. A. Goldsmith, J. Riggle, & A. Yu (Eds.), *Blackwell handbooks in linguistics. The handbook of phonological theory* (2nd ed., pp. 435–484). Malden, MA: Wiley-Blackwell.
- Shattuck-Hufnagel, S., & Turk, A. E. (1996). A prosody tutorial for investigators of auditory sentence processing. *Journal of Psycholinguistic Research*, 25(2), 193–247. <https://doi.org/10.1007/BF01708572>
- Sorensen, J. M., Cooper, W. E., & Paccia, J. M. (1978). Speech timing of grammatical categories. *Cognition*, 6(2), 135–153. [https://doi.org/10.1016/0010-0277\(78\)90019-7](https://doi.org/10.1016/0010-0277(78)90019-7)
- Turk, A., & Shattuck-Hufnagel, S. (2014). Timing in talking: What is it used for, and how is it controlled? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1658), 20130395. <https://doi.org/10.1098/rstb.2013.0395>
- Turk, A. (2010). Does prosodic constituency signal relative predictability? A Smooth Signal Redundancy hypothesis. *Laboratory Phonology*, 1(2), 11. <https://doi.org/10.1515/labphon.2010.012>
- Turk, A. E., & White, L. (1999). Structural influences on accentual lengthening in English. *Journal of Phonetics*, 27(2), 171–206. <https://doi.org/10.1006/jpho.1999.0093>
- Veilleux, N., Shattuck-Hufnagel, S., & Brugos, A. (2006). Transcribing Prosodic Structure of Spoken Utterances with ToBI. Retrieved from <https://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-911-transcribing-prosodic-structure-of-spoken-utterances-with-tobi-january-iap-2006>
- Venables, W. N., & Ripley, B. D. (2010). *Modern Applied Statistics with S* (4. ed.). *Statistics and computing*. New York: Springer.

- Wang, C., & Xu, Y. (2017). Effects of part of speech: Primitive or derived from word frequency? *Proceedings of 8th ExLing 2017, 19-22 June, Heraklion, Crete, Greece*, 113–116.
- Watson, D., Breen, M., & Gibson, E. (2006). The role of syntactic obligatoriness in the production of intonational boundaries. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(5), 1045–1056. <https://doi.org/10.1037/0278-7393.32.5.1045>
- Watson, D., & Gibson, E. (2004). The relationship between intonational phrasing and syntactic structure in language production. *Language and Cognitive Processes*, 19(6), 713–755. <https://doi.org/10.1080/01690960444000070>
- Watson, D., & Gibson, E. (2005). Intonational phrasing and constituency in language production and comprehension. *Studia Linguistica*, 59(2-3), 279–300. <https://doi.org/10.1111/j.1467-9582.2005.00130.x>
- Wightman, C. W., Shattuck-Hufnagel, S., Ostendorf, M., & Price, P. J. (1992). Segmental durations in the vicinity of prosodic phrase boundaries. *The Journal of the Acoustical Society of America*, 91(3), 1707. <https://doi.org/10.1121/1.402450>
- Yao, Y. (2011). The Effects of Phonological Neighborhoods on Pronunciation Variation in Conversational Speech (PhD thesis). UC Berkeley, Berkeley.