How morphological structure affects phonetic realization in English compound nouns

Melanie J. Bell, Sonia Ben Hedia & Ingo Plag

January 31, 2019

Abstract

Studies have shown that syntagmatic and paradigmatic aspects of morphological structure may impact on the phonetic realization of complex words (e.g. Sproat and Fujimura 1993; Kuperman et al. 2007; Smith et al. 2012; Schuppler et al. 2012; Lee-Kim et al. 2013; Cohen 2014a; Zimmermann 2016; Plag and Balling 2017; Lõo et al. 2018, among many others). The majority of these studies have been concerned with affixes, often focusing on the acoustic properties of the segments at a morphological boundary. Similar studies of compounds are still rare (but see Kuperman et al. 2007 on paradigmatic effects on duration and Kunter and Plag 2016 on effects of constituent structure in triconstituent compounds on duration). This study explores the extent to which consonant duration at compound internal boundaries in English is dependent on morphological strucure. Three competing hypotheses about the relationship between fine phonetic detail and morphological structure are tested. According to the Segmentability Hypothesis, higher morphological segmentability, i.e. a stronger morphological boundary, leads to lengthening (Hay 2004; Ben Hedia and Plag 2017; Plag and Ben Hedia 2018). The Informativity Hypothesis states that higher informativity leads to acoustic lengthening (Jurafsky et al. 2001; van Son and Pols 2003). Finally, the Paradigmatic Support Hypothesis says that stronger paradigmatic support leads to lengthening (Kuperman et al. 2007; Cohen 2014a, Tomaschek et al. 2018). To test these hypotheses, an experimental study was carried out with 62 compounds (types), taken from the BNC. The compounds were spoken by 30 speakers, yielding more than 1500 acoustic tokens overall. The data provide no support for the Segmentability Hypothesis, and only limited support for the Informativity Hypothesis. In contrast, the Paradigmatic Support Hypothesis made the correct prediction: N1 family size inversely correlated with consonant duration at the compound-internal boundary, which means that less paradigmatic uncertainty leads to longer durations. These results are in line with the idea that segment duration is lengthened under higher functional load, but shortened under functional uncertainty (Wedel et al. 2013; Tomaschek et al. 2018; Tucker et al. 2018).

1 Introduction

Many studies have shown that syntagmatic and paradigmatic aspects of morphological structure may impact on the phonetic realization of complex words (e.g. Sproat and

Fujimura 1993; Kuperman et al. 2007; Smith et al. 2012; Schuppler et al. 2012; Lee-Kim et al. 2013; Cohen 2014a; Zimmermann 2016; Plag and Balling 2017; Lõo et al. 2018, among many others). By 'syntagmatic' we mean the relationship between elements that occur in linear order in a stretch of speech or writing, while by 'paradigmatic' we mean the relationship of a given element to elements in absentia. This notion of 'paradigm' covers the classical inflectional paradigm, but also other kinds of sets of words that are morphologically related, including morphological categories, such as all words with the suffix *-ness*, and morphological families, such as all derived words containing a certain base, or all compounds that share a particular left or right constituent.

The majority of such studies have been concerned with inflectional and derivational affixes, often focusing on the acoustic properties of the segments at a morphological boundary (e.g. Smith et al. 2012; Lee-Kim et al. 2013). Investigations of morphologically induced phonetic variation in compounds are still rare but studies like Kuperman et al. (2007) or Kunter and Plag (2016) suggest that these types of complex words show similar effects. The present study extends this line of research by investigating the question of how consonant duration at compound-internal boundaries in English is dependent on morphological structure.

Insights into the relationship between morphological structure and phonetic implementation have important implications for theories of the mental lexicon and of speech production, perception and comprehension. Strictly feed-forward models of speech production (such as Levelt et al. 1999) and theoretical models of morphology-phonology interaction (e.g. Kiparsky 1982; Bermúdez-Otero 2018) rely on the distinction of lexical vs. post-lexical phonology and phonetics, excluding the possibility that morphological information influences the phonetic realization since this information is not available at the articulation stage.These theories are therefore incompatible with the findings mentioned above.

There have been some attempts to explain the unexpected phonetic effects of morphological structure. In particular, there are three published hypotheses that try to reconcile the unexpected phonetic findings with established ideas in various fields. First, according to the Segmentability Hypothesis (originating in the work of Hay 2003) the strength of a morphological boundary has an effect of phonetic implementation: higher morphological segmentability, i.e. a stronger morphological boundary, leads to lengthening (Hay 2007; Ben Hedia and Plag 2017; Plag and Ben Hedia 2018). Second, the Informativity Hypothesis states that a higher information load of some linguistic unit in speech leads to acoustic lengthening (e.g. Jurafsky et al. 2001; van Son and Pols 2003). This has been shown for different kinds of units, including morphological units (see Hanique and Ernestus (2012) for discussion). Finally, the Paradigmatic Signal Enhancement Hypothesis takes the structure of the morphological paradigm as its starting point and says that stronger paradigmatic support leads to acoustic lengthening (Kuperman et al., 2007; Cohen, 2014a; Tomaschek et al., 2018).

In this paper we test the three hypotheses by studying the duration of consonants at compound-internal boundaries in English. An experimental study was carried out with 62 compounds (types), taken from the BNC, spoken by 30 speakers, yielding more than 1500 acoustic tokens overall. The data provided no support for the Segmentability Hypothesis, and only limited support for the Informativity Hypothesis. In contrast, the Paradigmatic Signal Enhancement Hypothesis made the correct prediction: the family size of the first noun inversely correlated with consonant duration at the compound-internal boundary, which means that smaller paradigmatic uncertainty leads to longer durations. This result is in line with research that has shown that segment duration is lengthened under higher functional load, but shortened under functional uncertainty (Wedel et al. 2013; Tomaschek et al. 2018; Tucker et al. 2018).

The paper is structured as follows. In the next section we discuss in more detail the theoretical underpinnings of this study, developing the hypotheses to be tested. Section 3 introduces our methodology, which is followed by the presentation of our results in section 4. Section 5 discusses the theoretical implications of our findings.

2 Morphological structure and phonetic realization

As mentioned in the introduction, recent research on morphologically complex words has found evidence for correlates of morphological structure in the speech signal, i.e. in the way complex words are pronounced. Most such studies have focused on durational properties, but other aspects have also been investigated (e.g vowel formants, center of gravity, velarization). This line of research is important because it puts to test theories in different areas of linguistics: morphological theory (i.e. theories of phonology-morphology interaction in particular), theories of the mental lexicon (i.e. the representation and processing of complex words), and theories of speech production and perception. For all these theories there is a rather difficult problem to solve: How can morphological properties (e.g. the size of an inflectional paradigm or the strength of a suffix boundary) influence articulation in such a way that these properties have reflexes in the acoustic make-up of complex words? While the details of a solution to this problem are still largely unclear (and will probably be out of reach for some time to come), there are at least three approaches that have something to say about a possible relationship between morphological make-up and phonetic detail: Morphological segmentability, informativity, and paradigm structure.[1] We will discuss each in turn.

2.1 Morphological Segmentability

It is a general and well-established assumption across theoretical camps that there are weaker and stronger morphological boundaries. The strength of boundaries is usually diagnosed by a syndrome of structural, semantic and phonological properties. Weaker morphological boundaries are associated with lower productivity of the category in question, more bound bases, greater semantic opacity and enhanced phonological integration. Crucially, at the phonological level, the words with weaker boundaries show morpho-phonological alternations such as stress shift, resyllabification or assimilation. One theory that attempts to account for these phenomena is Lexical Phonology (e.g. Kiparsky 1982; Bermúdez-Otero 2018), where different lexical strata are posited to account for observable differences in boundary strength. In other words, boundary strength is taken to be categorical: level 1 or level 2.

[1] There is yet a fourth approach, prosodic phonology. In this approach different types of prosodic boundaries may correlate with different phonetic properties. Since the present paper deals with only one type of prosodic structure, i.e. compounds, we do not test this approach, and refrain from a discussion.

Assuming differences in boundary strength is in line with dual route models of morphological processing, i.e. with models that allow both whole-word storage and morphological segmentation. Hay (2003) argues that words with a strong boundary are more likely to be segmented and their constituent morphemes processed individually, while words with a weak boundary are more likely to be processed holistically. In contrast to Lexical Phonology, Hay's approach takes boundary strength to be gradient, influenced by parameters such as semantic transparency, phonological transparency, and relative frequency of the the complex word and its base. Phonetically, words with weaker boundaries are expected to show more phonetic reduction across the morpheme boundary than words that have a strong boundary. For example, in contrast to less frequent and easily segmentable derived words, such as *imprison-ment* or *compact-ly*, high frequency words like *govern-ment* or *exact-ly* show stronger reduction effects, such as the loss of the second syllable in *government*, or of the /t/ in *exactly* (cf. Hay 2003).

A number of empirical studies have found that morphological segmentability systematically affects the acoustic realization of complex words. Sproat and Fujimura (1993) and Lee-Kim et al. (2013), for example, show that the realization of /l/ in English depends on the strength of the morphological boundary it occurs at. Stronger boundaries go together with longer duration and stronger velarization of /l/. In a study of English *-ly*-suffixed words, Hay (2003) found less acoustic reduction with more segmentable derivatives than with less segmentable derivatives. In Hay (2007), she found similar results for the prefix *un-*, i.e. the prefix was less reduced in more segmentable derivatives than in less segmentable derivatives. Similarly, Ben Hedia and Plag (2017) show that a higher degree of segmentability correlates with less reduction in prefixed words. In their study on three English prefixes, locative *in-*, negative *in-* and *un-*, they found that the least segmentable prefix, locative *in-*, features the shortest nasal and the most segmentable prefix, *un-*, features the longest nasal. Smith et al. (2012) investigated prefixed words with *dis-* and *mis-*. In both categories, there are highly segmentable words (such as *mistime*, *mistype*, *displeased*, *discolored*), called 'prefixed' by these authors, and less easily segmentable ones, (e.g. *mistake*, *discovered*, *distorted*), called 'pseudo-prefixed'. The analysis of different phonetic characteristics (duration, formant structures, amplitude, spectral moments) shows that the prefixes in the pseudo-prefixed words have shorter durations than in the prefixed words, and that segments straddling a weaker morphological boundary show phonetic characteristics that are closer to those of morpheme-internal sequences of the same type. That the segmentability of affixes affects their phonetic realization is also shown in Plag and Ben Hedia (2018), who find that the duration of the two English prefixes *un-* and *dis-* is affected by their segmentability. A higher degree of segmentability goes together with longer durations.

However, while there is evidence that a higher degree of segmentability leads to less phonetic reduction in complex words, the scope and nature of these effects is not quite clear. In Plag and Ben Hedia (2018), the effect was not observed for all of the investigated affixes, since the duration of *in-* and *-ly* was not affected by any of the tested segmentability measures, and in Hay (2007), the effect was not observed for all of the tested speakers. Furthermore, a number of studies failed to find any segmentability effect on the acoustics of complex words (e.g. Schuppler et al. 2012, Bürki et al. 2011).

As discussed in the overviews by Hanique and Ernestus (2012) and Ben Hedia (2018), the deviating findings across studies might be caused by the application of different segmentability measures, as well as by differences in the structure investigated. Thus, some studies looked at suffixed words (e.g. Hay 2003; Schuppler et al. 2012), others at prefixed words (Ben Hedia and Plag 2017), some looked at effects at the morpheme level (Hay 2007; Plag and Ben Hedia 2018), others at the segment level (Hay 2003; Bürki et al. 2011; Ben Hedia and Plag 2017), some looked at pre-boundary reduction (Hay 2003; Ben Hedia and Plag 2017), others at post-boundary reduction (Schuppler et al. 2012). Some studies investigated inflection (e.g. Bürki et al. 2011; Schuppler et al. 2012), others derivation (Hay 2007; Ben Hedia and Plag 2017). There are no studies available yet that have tested phonetic effects of segmentability on compounds.

Although there is no empirical work specifically investigating the effect of segmentability on the acoustics of compounds, there are a number of studies which provide evidence that the semantic transparency of a compound, which is some kind of segmentability measure, affects the way it is processed. For example, Ji et al. (2011) found that when meaning decomposition is encouraged, transparent compounds are processed faster than opaque compounds. For Dutch, Zwitserlood (2007) found that, unlike semantically transparent compounds, semantically opaque compounds do not prime the associates of their constituents. According to the author, these results indicate that on a semantic level, transparent compounds are linked to the their constituents while opaque compounds are not. In a similar vein, MacGregor and Shtyrov (2013) argue that semantically transparent compounds are processed via their parts while semantically opaque compounds are processed as a whole.

One could assume that the alleged differences in processing between semantically transparent and semantically opaque compounds are mirrored in their articulation. Transparent, more segmentable compounds, which are processed via their parts, would show less acoustic reduction than opaque, less segmentable compounds, which are processed as a whole. This assumption fits in well with the segmentability effects found on the acoustics of derived words, where a higher degree of segmentability leads to less acoustic reduction. Some indirect indication in this direction may come from a study by Kunter and Plag (2016). In their study of triconstituent English compounds (e.g. [[*day care*] *center*], the authors investigated whether the internal bracketing of compounds affects the acoustic realization of the constituents. It was found that the durational properties of the constituents straddling the boundary at the immediate constituent level are indicative of the hierarchical structure of the compound, and of the strength of the boundary between the immediate consitituents. The weaker the internal boundary between the two words of the complex constituent, the longer the duration of the third constituent, and the shorter the duration of the embedded constituent next to it.

In the present study we will test the idea that factors facilitating morphological segmentation lead to phonetically longer pronunciations, using English compounds as our data. We will focus our attention on what happens at the internal boundary of the compound and consider the duration of the consonants at this boundary. We will test the 'Segmentability Hypothesis', which we specify for our purposes as in (1):

(1) **Segmentability Hypothesis**
The more segmentable a compound, the longer is the duration of consonants at compound-internal boundaries.

2.2 Informativity

Many studies have shown that the information load (henceforth 'informativity') of a linguistic unit affects its phonetic realization. Speakers pronounce words faster, i.e. with shorter duration, when they are contextually expected, and thus add little information. In other words, the more informative a unit is, the less reduction one finds. This has been shown for different types of unit: individual segments (e.g. van Son and Pols 2003; van Son and van Santen 2005), syllables (Aylett and Turk 2006), and words (e.g. Jurafsky et al. 2001; Bell et al. 2009; Seyfarth 2014).

There are also studies that looked at the morpheme as the crucial unit. For example, in a study on Spanish word-final /s/, Torreira and Ernestus (2012) found that /s/ suffixes in redundant morphosyntactic contexts were more likely to reduce than other word-final /s/ segments. For instance, the /s/ in *cuatro cosa**s*** 'four things' is shorter than the /s/ in *quiero cosa**s*** 'I want things'. Working on English final /s/, Cohen (2014a) finds a similar effect, in that third person singular *-s* is pronounced with shorter duration if it is contextually more probable.

While the studies just mentioned looked at the syntagmatic dimension of informativity, i.e. at how informative a string may be in its syntagmatic context, there is also some work that has looked at informativity on the paradigmatic axis. In their study of Dutch complex words ending in *-igheid*, Pluymaekers et al. (2010) measured informativity indirectly in terms of two categories (*ADJ-igheid* vs. *[ROOT-ig]-heid*) that have high and low paradigmatic informativity at the /xh/ transition (i.e. the transition between *-ig-* and *-heid*), respectively. In that data set /xh/ was shorter for the *[ROOT-ig]-heid* words, which means that the category with low informativity correlated with shorter durations.

The aforementioned studies all use probabilistic measures of informativity in the spirit of Information Theory (Shannon, 1948). Using both probabilistic and semantic measures of informativity, Ben Hedia (2018) also argues for an effect of morphological informativity on the acoustics of complex words. She claims that the acoustic realization of morphological geminates in English affixed words (e.g. /nn/ in *unnatural*) depends on the informativity of the affix. She finds that the more informative an affix is, the longer is the duration of the double consonant.

To sum up, there is compelling evidence for the assumption that more informative morphemes are less reduced than less informative morphemes. Even though there is no empirical work on the influence of informativity on the acoustics of compounds, one can nevertheless assume that similar effects can be found for this type of complex word. For our study, this leads to the 'Informativity Hypothesis' spelled out in (2):

(2) **Informativity Hypothesis**
The more informative the constituents of a compound, the longer is the duration of consonants at compound-internal boundaries.

2.3 Paradigm Structure

There is ample evidence that paradigmatic structure plays an important role in the processing of inflected words, derived words and compounds (see, for example, Baayen et al. (1997); Milin et al. (2009a) on inflection, Schreuder and Baayen (1997); Milin et al. (2009b); Kuperman et al. (2009) on derivation, van Jaarsveld et al. (1994); Kuperman et al. (2008, 2010) on compounds). It is therefore not far-fetched to think that it may also affect speech production and lead to paradigm-specific acoustic patterns. And indeed, such evidence has been found.

The effects of paradigmatic structure on processing as well as on speaking are usually measured using numerical predictors gleaned from a word's paradigm. Such measures can be the size of the paradigm, the number of certain competitors in a paradigm, the probability of a given form in a paradigm, or the paradigm entropy. These measures are thus related to the informativity measures discussed in the previous subsection, but they measure aspects of the paradigmatic, not of the syntagmatic, structure, i.e. paradigm-specific informativity.

Several studies have found that paradigmatic structure impacts on pronunciation, and two opposite effects have actually been reported: enhancement and reduction. Let us look at enhancement first. 'Enhancement' refers to effects in which more paradigmatic support (e.g. higher paradigmatic probability) leads to longer durations or more distinct pronunciations. Interestingly, this effect works in the opposite direction from the reduction effects caused by syntagmatic informativity.

Kuperman et al. (2007), for example, investigated the interfixes *-s-* and *-en-* in Dutch compounds and found that the more probable an interfix is, the longer is its duration. Cohen (2014a) not only found the syntagmatic effect described in the previous subsection, but also discovered a paradigmatic enhancement effect with increasing relative paradigmatic frequency (the frequency of third person singular form in relation to the frequency of the plural form). The more probable the third person singular form in its verbal paradigm, the longer the suffix. Investigating the vowel formants in Russian verbs, Cohen (2014b) demonstrates that with rising paradigmatic probability of the verb form in question, the vowels are pronounced more distinctly. For Estonian case-inflected nouns, Lõo et al. (2018) found that inflectional paradigm size not only predicts word naming latencies, but also acoustic durations. Smaller inflectional paradigms and smaller morphological families, i.e. higher paradigmatic probability (lower paradigmatic informativitity), go together with longer durations. In a recent study of final *-s* in English, Tomaschek et al. (2018) use a Naive Discriminative Learning network to predict the duration of different types of final *-s* (i.e. non-morphemic, plural, genitive, genitive plural, cliticized auxiliaries *has, is*). The results indicate that when a final /s/ provides strong support for the targeted inflectional category the /s/ is articulated with longer duration. When /s/ is not a good discriminative cue, i.e. it creates uncertainty about its morphological function by providing support for different functions, its duration is decreased by the speaker.

For the present study we adopt the idea of paradigmatic enhancement and test the hypothesis in (3):

(3) **Paradigmatic Support Hypothesis**
The more paradigmatic support, the longer the duration of consonants at compound-internal boundaries.

3 Methodology

3.1 Data

We are looking at the consonants at compound boundaries, as, for example, in:

(4) *stea**m** engine*

(5) *crea**m** **m**ini*

(6) *survey **m**anager*

The consonant is either part of the first noun ('N', as in (4)), the second noun ('N2', as in (6)), or of both (as in (5)), which gives us the possibility of testing which factors affect which part of the boundary. In other words, if there is reduction, does it take place before the boundary, after the boundary or at both sides of the boundary?

We especially wanted to include compounds such as *crea**m** **m**ini*, with a double consonant at the boundary, to maximise our chances of finding a paradigmatic enhancement effect. The only previous report of such an effect for compounds is that reported by Kuperman et al. (2007), who found a paradigmatic enhancement effect on the duration of interfixes in Dutch compounds. Although English does not have interfixes, we reasoned that we might see a similar effect on the segments at compound-internal boundaries, perhaps especially on morphological geminates, since in such cases a single articulatory gesture straddles the boundary. Geminates may therefore be subject to influence by the lexical properties of both constituents, just as interfixes are. In the present study we focus on the consonants /m/, /n/ and /s/, since it has been shown (e.g. by Ben Hedia 2018) that these sounds may show morphological gemination in English.

English compounds show considerable variation in orthographic representation between spaced, hyphenated and unspaced spellings. However, unspaced and hyphenated spellings tend to correlate with high frequency and lexicalisation. In order to find a sample of attested compounds with a wide range of frequencies, and at the same time avoid the complicating factor of varied spelling, we therefore decided to focus exclusively on spaced compounds.

The compounds used in the present study were selected from the spoken section of the British National Corpus (John Coleman et al., 2012). Using only the spoken section of the corpus ensures that the resulting compounds have been spontaneously produced by a speaker at least once. The Lancaster interface was used to search for strings of two nouns, excluding strings that crossed a sentence boundary or that included a pause or any other form of interruption, e.g. a cough, between the two nouns. The corpus queries also specified that the word after the second noun should not be another noun, an adjective or a possessive. This restricted the searches to strings of exactly two nouns and excluded combinations that were part of a larger compound construction. Despite these precautions, the strings were subsequently checked in context to ensure that they did represent constructions in which the first noun modified the second. We take the view, following e.g. Bauer (1998); Plag et al. (2008); Bell (2011), that all such constructions can be classed as compounds. At this stage, types in which the two nouns were identical or either noun was hyphenated, as well as proper names, appositive constructions and vocatives, were also excluded from the data.

Phonological transcriptions of the constituent nouns of the compounds were extracted from CELEX and in cases where a word did not appear in CELEX they were supplemented by manual transcription. These transcriptions were then used to identify types in which the first word ended with one of the consonants /s/, /m/ or /n/, and the second word started with the same phoneme. From this set, we selected only those combinations in which neither the word final consonant nor the word initial consonant formed part of a cluster. To be able to compare double consonants with both word final and word initial single consonants, we also used the transcriptions to select compounds in which either the first constituent ended with /s/, /m/ or /n/ and the second word began with a vowel, or the second constituent started with one of these consonants and the first word ended with a vowel. Again, we excluded types with word initial or word final clusters.

We further restricted ourselves to types in which, according to CELEX or our manual transcriptions, the lexical stress of the second noun fell on the first syllable of that noun. In all the types selected, there was therefore either a single or a double consonant at a compound boundary and between two vowels, the second of which bore lexical stress (though this was not necessarily the main stressed syllable of the whole compound, only of the second constituent). Some examples are shown in Table 1.

Table 1: Examples of experimental items

m			
m#V	*calcium intake*	*problem area*	*bum ache*
m#m	*cream mini*	*pandemonium model*	*claim money*
V#m	*company money*	*media men*	*polo mint*
n			
n#V	*kitchen area*	*pavilion end*	*swan inn*
n#n	*pen knife*	*woman novelist*	*criterion number*
V#n	*polo neck*	*key note*	*mercury number*
s			
s#V	*surface area*	*bonus element*	*police act*
s#s	*dress sense*	*gas side*	*peace settlement*
V#s	*survey site*	*eye sore*	*tuna sandwich*

From the set of compounds described in the previous paragraph, we selected a subset to use in our study. In selecting the subset we aimed to achieve as wide and balanced a range as possible across the following criteria:

- number of syllables in first noun
- number of syllables in second noun
- weight of final syllable in first noun (i.e. strong or weak)
- expected position of compound stress (i.e. on first or second noun)

- vowel phoneme preceding the consonant(s)
- vowel phoneme following the consonant(s)

In other words, items were selected to enhance the diversity of the data with respect to these criteria, and avoid a bias towards any specific syllable structure, stress or vowels. On the other hand, items were excluded if they were unique in terms of any of these variables, since that would have introduced a confound between compound and condition. In cases where more than one compound satisfied all these constraints, the final selection was made randomly. These procedures resulted in a list of compounds with 19 compound types with /m/ at the boundary, 19 types for /n/, and 24 types for /s/ (see Table 3 below for a more detailed overview).

3.2 Experimental set-up

Spoken tokens of all the compounds in our final dataset were elicited from 30 native speakers of British English, who read the compounds presented in carrier sentences on a computer screen. Each compound was embedded in two different carrier sentences:

(7) *They talked about the [compound] again.*

(8) *She told me about the [compound].*

The two sentences differ with respect to whether the compound occurs in final position, to allow for any lengthening or shortening effects of phrasal position to be included in the analysis. Each participant read each compound only once, either in sentence (7) or in sentence (8). However, overall each participant saw an equal number of both sentence types, and each compound was included in an equal number of tokens of each sentence type. The sentences were mixed with an equal number of unrelated filler sentences, which were the experimental items for another study. Since the filler sentences had a variety of different structures, they served to break up the repetitiveness of our carrier sentences and reduce the risk of a list-like intonation developing. Each participant saw the items, including fillers, in a different randomised order.

Each sentence was presented on two consecutive slides. The first slide of each pair asked the participant to read the sentence silently, while the second slide instructed them to read the sentence aloud. The silent reading phase was intended both to encourage semantic processing of the sentence and to reduce the risk of performance errors in the subsequent reading aloud. There was an initial training phase, and participants could move through the presentation at their own pace.

The recordings were produced in a sound-proof booth, digitised at 44.1kHz using a Tascam HD-P2 digital recorder and a Senheiser ME 64 cardioid microphone, with participants seated 15cm from the microphone and recording levels set for each participant.

3.3 Acoustic measurements

After recording the sentences, we manually segmented the data and transcribed them phonetically using the software Praat (Boersma and Weenink 2014). We annotated the segments in question, as well as the preceding and the following segment. The annotation for *steam engine*, for example, included the segmentation of [i], [m] and [ə].

The segmentation was carried out according to criteria that relied on the visual inspection of the waveforms and spectrograms of the items. These criteria were based on the segmentation criteria applied in Ben Hedia (2018), which in turn were based on the features of specific sounds as described in the phonetic literature (e.g. Ladefoged 2003). As all of the investigated consonants occur in intervocalic position, we concentrated on the differences between the pertinent consonants and vowels. Nasals have a regular waveform which has a lower amplitude than the waveform of vowels. Furthermore, their formants are quite faint in comparison to those of vowels. This can be seen in figure 6, which shows a sample segmentation of the word *steam engine*. Fricatives have, in contrast to vowels, an aperiodic waveform and are therefore quite easy to identify in intervocalic position. All boundaries were set at the nearest zero crossing of the waveform.

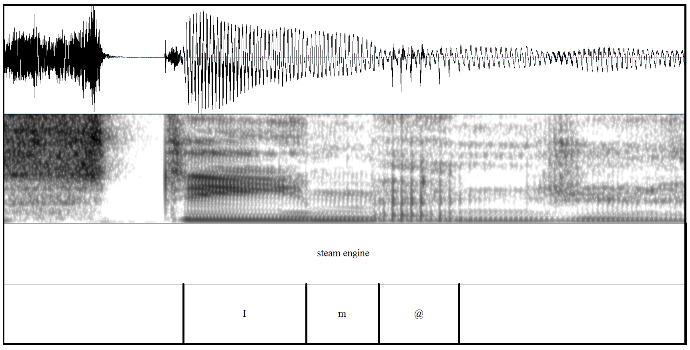

Figure 1: Annotation of the compound *steam engine*

Double consonants (e.g. /mm/ in *cream mini*) were treated as one segment in the annotation when no boundary between the two identical consonants was discernible. If there was a visible boundary between the two consonants, both consonants were segmented. This was the case when the speaker produced a pause between the first and the second constituent. Such tokens were subsequently excluded from the analysis.

The reliability of the segmentation criteria was verified by a set of trial segmentations. In these trials, three annotators used the criteria to segment the same 20 items. If there was any larger discrepancy in the placement of the boundaries, i.e. if any boundary differed from another by more than 10 milliseconds, the annotators discussed the discrepancy and refined the criteria in order to reduce the amount of intra-annotator variation. These trial segmentations were repeated until all boundaries were reliably placed with only small variations. For the final measurement, each annotator worked on a disjunct set of items. After the segmentation process was completed, a script was used to measure and extract word duration, constituent durations, the duration of the consonants in question, as well as the duration of its preceding and following segment in milliseconds.

3.4 Predictor variables

To test the three hypotheses under consideration, we extracted a number of frequency-based measures from ukWaC (https://www.webarchive.org.uk/ukwa/), a corpus of more than 2 billion words from the .uk internet domain. These included:

- SpellingRatio: (f(concatenated)+f(hyphenated))/f(spaced)

- CompoundFrequency: N2 lemmatised; all spelling variants
- Constituent frequencies: N1Frequency lemmatised, N2Frequency lemmatised; based on all spelling variants
- Conditional frequency N2 given N1, CondFreqN2: f(compound)/f(N1)
- Constituent family sizes: N1FamilySize and N2FamilySize, both based on spaced NN strings occurring within a sentence
- Entropy (N1Entropy, N2Entropy): The entropy of the N1 constituent family and of the N2 constituent family, using token frequencies of the different compounds in each family

Frequency measures and family sizes were log-transformed before entering them into the statistical analysis. Let us now see how these measures relate to the three hypotheses.

3.4.1 Segmentability

The segmentability of compounds can be tested on several grounds. In this study we use two measures as correlates of segmentability: spelling ratio and N1 family size. Spelling ratio refers to whether the compound is spelled with or without a space. It is thus the ratio of the frequency with which the compound is spelled without a space to the frequency with which the compound is spelled with a space. It is assumed that the space indicates a higher segmentability. Thus, a compound with more spaced spellings is more segmentable than a compound with more non-spaced spellings, and it should thus feature longer consonants at its internal boundary. In other words, according to this hypothesis, spelling ratio should be inversely correlated with duration: higher spelling ratio means less decomposable, leading to shorter durations.

The second measure of segmentability is N1 family size. It is assumed that the larger the N1 family, the more productive is N1 as a compound modifier, and greater productivity has been show to be associated with greater decomposability of complex words (cf. Hay and Baayen 2003).

3.4.2 Informativity

Informativity is highly related to concepts of probability and predictability. A morpheme that is less probable is less predictable, and in turn more informative. A morpheme that is highly probable is very predictable, and thus less informative. It is thus reasonable to use predictability measures as proxies for informativity.

We will test six different types of predictability: compound frequency, constituent frequencies, conditional probability of the consonant in question, conditional probability of N2, family size of N2 and the entropy of the N2 family. While compound frequency refers to the frequency of the whole compound, constituent frequencies refer to the frequencies of the two constituents of the compound. In both cases, it is predicted that a higher frequency should go together with shorter durations as it indicates high probability, i.e. low informativity.

We calculated the conditional probability at the segment level by calculating the probability of the consonant at the compound internal boundary given the preceding vowel. This measure was based on transitions within monomorphemic words in CELEX.

The conditional probability of N2 refers to the probability N2 given N1, e.g. the probability of *manager* given the word *survey* for the compound *survey manager*. The higher the conditional probability of N2, the less informative is N2, and the shorter the consonant is expected to be. The size of the constituent family of N2 may have a similar effect. In compounds with a large N2 family, N2 is highly predictable, hence rather uninformative and therefore prone to reduction. In other words, increasing N2 family size reduces the informativity of N2, which should lead to shorter durations. The same effect should also be observable with the entropy of the N2 family. Entropy has been established as an alternative measure of paradigmatic complexity (e.g. Milin et al. 2009c; Kuperman et al. 2007). A high entropy of the N2 family means that N2 is not informative, which leads to phonetic reduction at the boundary, as shown in Kuperman et al. 2007 for Dutch compounds.

3.4.3 Paradigmatic Support

For compounds, one can use N1 family size as a measure for paradigm size. The higher the N1 family size, i.e. the more possible values there are for N2, the more spread is the activation, and the less support there is for a particular N2, given N1. Given the paradigmatic enhancement hypothesis, it is predicted that an increased N1 family size leads to shorter consonant durations at the compound-internal boundary.

The role of N2 family size can be seen as the reverse of N1 family size. A larger N2 family size means that a particular N2 is more likely to occur with more different N1 constituents, and is therefore more certain to occur. According to the paradigmatic enhancement hypothesis this means that larger N2 family sizes would lead to longer durations.

N1 and N2 family sizes thus work in opposite directions. Increasing N1 family size increases insecurity for the transition from N1 to N2, while increasing N2 family size decreases insecurity for the transition from N1 to N2.

The entropy of the N1 family size should have similar effect on consonant duration as N1 family size. The higher the entropy of the N1 constituent family, the higher the uncertainty about what comes as the second constituent. The predictions concerning the respective effects of N1 family entropy are thus parallel to those of family size: higher N1 entropy should lead to shorter durations at the boundary (cf. Kuperman et al. 2007).

Note that the segmentability hypothesis and the paradigmatic enhancement hypothesis make conflicting predictions about the effect of N1 family size, and this predictors can therefore be used to evaluate them against each other. Table 2 summarizes the variables of interest and the predictions concerning the effect of these variables, as derived from the three hypotheses.

3.4.4 Control variables

In addition to our predictors of interest, we also included a number of control variables in our models. These were:

Table 2: Summary of predictors and their predicted effects

Measure	**Segmentability Hypothesis**	**Informativity Hypothesis**	**Paradigmatic Support Hypothesis**
Higher constituent frequencies		shorter duration	
Higher compound frequency		shorter duration	
Higher conditional probability		shorter duration	
Higher spelling-ratio	shorter duration		
Greater N1 family size	longer duration		shorter duration
Higher N1 entropy			shorter duration
Greater N2 family size		shorter duration	
Higher N2 entropy		shorter duration	

Boundary type (C#C, C#V or V#C). We included BoundaryType for two reasons. The first is, that phonetic studies have shown that the duration of consonants may be influenced by the phonetic context in which they occur (e.g. Umeda 1977). The second reason is that the geminates may behave differently from the other consonants because they straddle the word boundary, and may therefore be especially prone to effects arising from the combination of the two constituents.

Consonant (/m/, /n/ or /s/). Consonant controls for the inherent duration differences between the three consonants.

Local speech rate. SpeechRate was computed by dividing the duration of the compound by the number of segments the compound contains. Obviously, a faster speech rate leads to shorter durations of individual segments.

Number of syllables in N1, and Number of Syllables in N2. It has been shown (e.g. by Lindblom 1963; Nooteboom 1972) that segments may tend to be shorter if the words in which they occur have more syllables. This effect can be conceptualized as a kind of compression effect, where words with more syllables undergo reduction. We therefore included syllable counts of the two constituents (N1SyllableCount and

N2SyllableCount) in our set of covariates.

Presentation order of items. The variable Order was included to control for effects of variability in attention or fatigue across the duration of the experiment.

3.5 Statistical analysis

We carried out mixed effects regression analysis using the `lme-4` package in R (Bates et al., 2014). The dependent variable was the duration of the consonant at the compound internal boundary, and before analysis we trimmed the data to remove outliers with very long or short durations. We also removed outliers with respect to local speech rate. This process resulted in a loss of 22 data points, about 1.4% of the data.

Many of our variables correlate with each other, which means that one needs to take care of potential collinearity issues. We adopted the following strategy. We first determined which variables were highly correlated. We then modeled the effect of each of these variables on consonant duration in individual separate models. Of two correlating variables we then included the variable that had the strongest effect on consonant duration and excluded the correlating variable from further analyses. The resulting set of variables was checked for remaining potential collinearity by using the collin.fnc function of the LanguageR package (Baayen, 2010).

We also included interactions based on two considerations. First, we expected that at least some of the variables would play out differently depending on the position of the consonant, i.e. our variable BoundaryType. Second, we expected the different consonants to potentially show differences in their slopes in relation to various other predictors. The initial models therefore three-way interactions between the variables BoundaryType and Consonant and all our variables of interest plus SpeechRate. None of the three-way interactions turned out to be significant, only some two-way interactions did.

In addition to the fixed effect predictors and control variables described above, our initial models also included random intercepts for participant, item, annotator and compound position (sentence-final or not). We also tested more complex random effects structure by including by-participant slopes for the continuous fixed predictors. A random effect was kept in the model if their inclusion led to a significantly better model according to a log-likelihood test. The final model was derived by stepwise elimination of non-significant predictors using the `step` function of the `lmerTest` package (Kuznetsova et al., 2017).

The number of types and tokens in the dataset for the final model is shown in Table 3.

Table 3: Distribution of types and tokens

m		
	Number of Types	Number of Tokens
m#V	5	117
m#m	9	222
V#m	5	125
Total	19	464

n		
	Number of Types	Number of Tokens
n#V	5	95
n#n	9	236
V#n	5	119
Total	19	450

s		
	Number of Types	Number of Tokens
s#V	5	96
s#s	14	373
V#s	5	118
Total	24	587

4 Results

The initial model was fitted as described in the previous section. The inspection of the residuals of the final model revealed an unsatisfactory distribution of the residuals. To address this problem, we Box-Cox transformed the consonant durations (λ=0.4242424) and, in a final step, removed data points with absolute residuals larger than 2.5 standard deviations, which resulted in the loss of 1.9 percent of observations. The resulting final model showed a normal distribution of the residuals (Shapiro-Wilk normality test, W=0.998, p=0.1207).

The final model includes random intercepts for Item and Participant. In addition there are five significant fixed effects, with four two-way interaction terms. The model is documented in Tables 4 and 5.

We will first discuss the effects of the control variables starting with the interactions involving the local speech rate.

Table 4: Fixed effects of the final model for duration of consonants at compound-internal boundaries

	Sum Sq	Mean Sq	NumDF	DenDF	F.value	Pr(>F)
BoundaryType	0.0607	0.0304	2	65.0587	63.6873	0.0000
Consonant	0.0042	0.0021	2	61.0814	4.4303	0.0160
SpeechRate	0.0631	0.0631	1	940.3013	132.3386	0.0000
N1FamilySize	0.0013	0.0013	1	51.9513	2.6571	0.1091
N2FamilySize	0.0076	0.0076	1	51.7836	16.0351	0.0002
BoundaryType:SpeechRate	0.0258	0.0129	2	1030.9250	27.0236	0.0000
BoundaryType:N1FamilySize	0.0119	0.0060	2	52.2104	12.5255	0.0000
Consonant:SpeechRate	0.0072	0.0036	2	1067.1454	7.5874	0.0005
Consonant:N1FamSize	0.0042	0.0021	2	51.8181	4.3690	0.0176

Table 5: Coefficients of the final model for duration of consonants at compound-internal boundaries

Random effects:					
Groups	Name	Variance	Std.Dev.	Corr	
Item	(Intercept)	0.0001219	0.01104		
Participant	(Intercept)	0.0002063	0.01436		
Residual		0.0004766	0.02183		
Number of obs: 1472, groups: Item, 62; participant, 30					
Fixed effects:					
	Estimate	Std. Error	df	t value	Pr(>|t|)
(Intercept)	0.4493	0.0320	63.5773	14.0452	0.0000
BoundaryType-C#C	0.2379	0.0250	65.2422	9.5236	0.0000
BoundaryType-C#V	-0.0027	0.0286	61.8274	-0.0961	0.9238
Consonant-m	-0.0914	0.0307	63.4604	-2.9731	0.0042
Consonant-n	-0.0691	0.0341	59.0136	-2.0288	0.0470
SpeechRate	-0.0045	0.0008	615.5444	-5.7568	0.0000
N1FamilySize	0.0042	0.0042	52.8396	1.0198	0.3125
21FamilySize	-0.0055	0.0014	51.7836	-4.0044	0.0002
BoundaryType-C#C:SpeechRate	-0.0045	0.0007	928.3848	-6.0794	0.0000
BoundaryType-C#V:SpeechRate	0.0004	0.0009	1037.1164	0.4110	0.6812
BoundaryType-C#C:N1FamilySize	-0.0154	0.0033	52.0772	-4.6139	0.0000
BoundaryType-C#V:N1FamilySize	-0.0034	0.0039	51.3655	-0.8625	0.3924
Consonant-m:SpeechRate	0.0016	0.0008	1039.6441	2.0560	0.0400
Consonant-n:SpeechRate	0.0029	0.0007	986.2001	3.8867	0.0001
Consonant-m:N1FamilySize	0.0040	0.0039	51.6084	1.0294	0.3081
Consonant-n:N1FamilySize	-0.0065	0.0046	51.8792	-1.4072	0.1653

The interactions involving SpeechRate are shown in Figures 2 and 3. They illustrate the effect of SpeechRate on consonant duration, modulated by Consonant (Figure 2) and by BoundaryType (Figure 3). Figure 2 shows three main things. First, as expected, the fricative is longer than either of the nasals, and /m/ is slightly longer than /n/. Second, as expected, SpeechRate has an influence on the consonant duration. Higher speech rate leads to shorter consonants. Third, this shortening effect is most pronounced for /s/ and least pronounced for /n/. Figure 3 also shows the general effect of speech rate on consonant duration. In addition, it also shows very clearly the effect of morphological gemination: double consonants are almost always longer then their singleton counterparts. Initial consonants in the second constituent are marginally longer than final consonants in the first constituent.

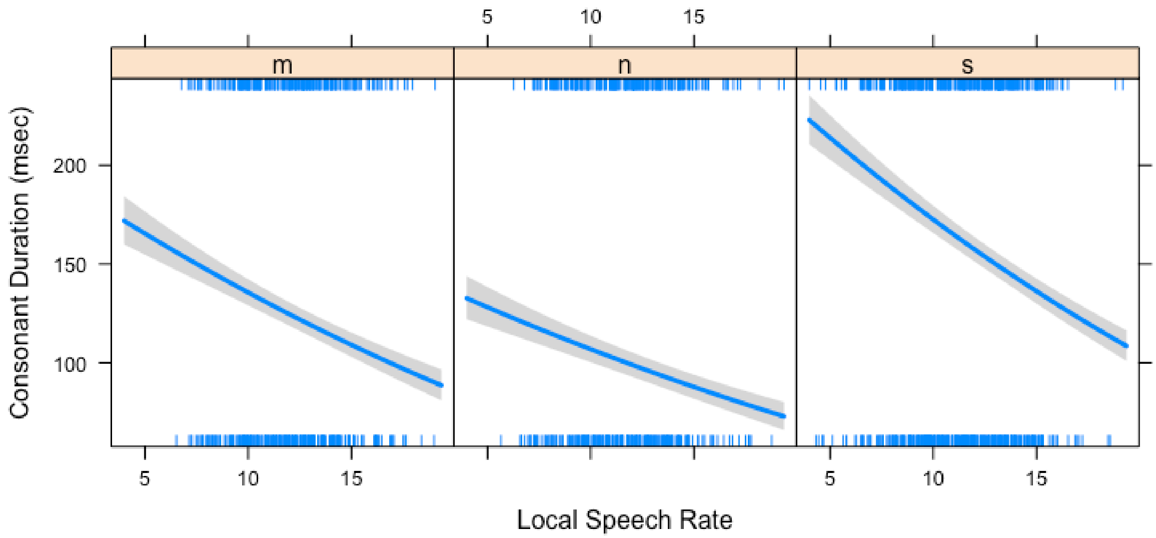

Figure 2: Partial effect of local speech rate by consonant in final model

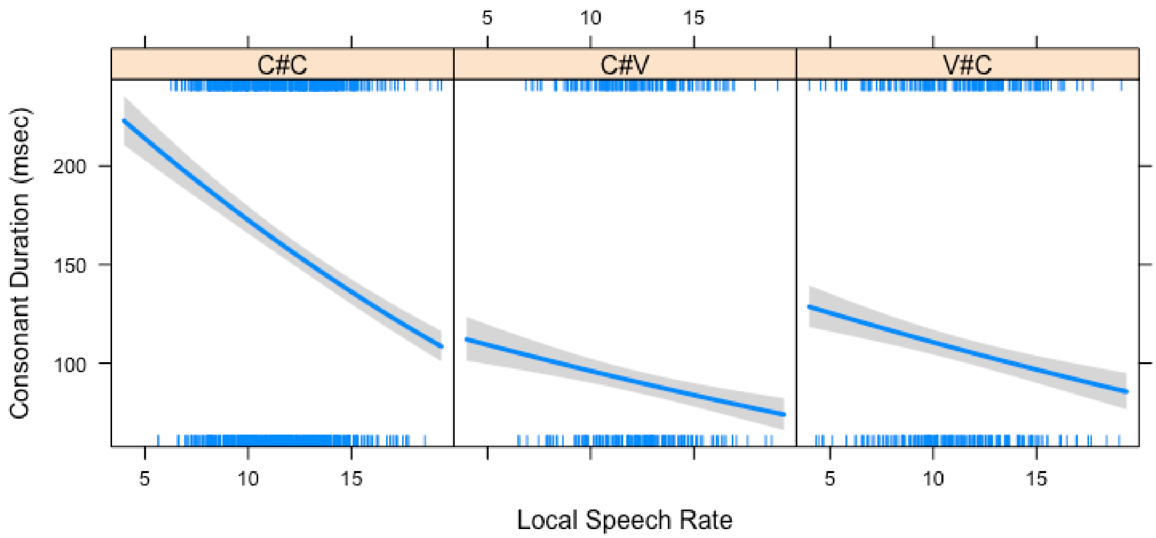

Figure 3: Partial effect of local speech rate by boundary type in final model

We turn now to the predictors of interest, that is to those variables which are of immediate relevance to our three hypotheses. The majority of these predictors did not reach significance. Neither compound frequency, nor spelling ratio, nor either of the two constituent frequencies survive in the final model. The conditional probability at the boundary also has no effect on consonant duration, irrespective of whether this is measured in terms of predictability of the second word or of the consonant given the preceding vowel.

The effects that we do find are for the constituent family sizes. These turned out to be better predictors of duration than the token-based entropy of the paradigms. Hence entropy measures are not in the final model. The effects of N1 family size are shown in Figures 4 and 5. Figure 4 shows the effect of N1 family size modulated by consonant. Again we see that the fricative is longer than the nasals and over most of the distribution of the data, /m/ is slightly longer than /n/. However, in all three cases duration falls slightly with increasing family size.

Remember that N1 family size is the predictor variable that most clearly enables us to differentiate between the segmentability hypothesis and the paradigmatic enhancement hypothesis. According to the segmentability hypothesis, duration would be expected to be positively correlated with N1 family size: the greater the family size of the modifier, the more productive the modifier, hence the more segmentable the compound, the stronger the boundary and the longer the consonant. This is not what we find. On the contrary, we find that consonant duration falls with increasing N1 family size as predicted by the paradigmatic enhancement hypothesis.

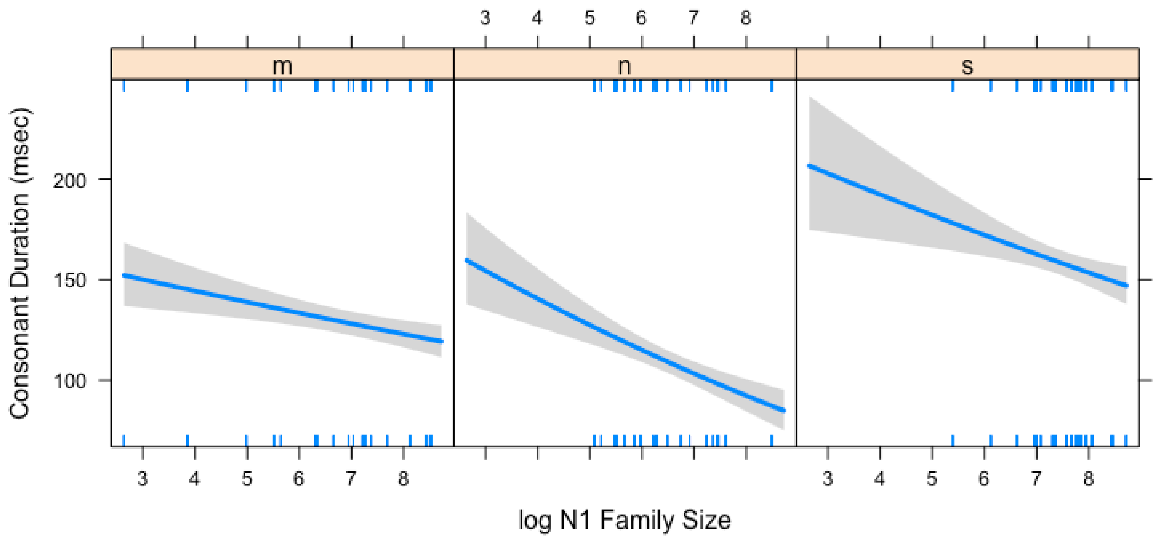

Figure 4: Partial effect of N1 family size by consonant in final model

Figure 5 shows the interaction of N1 family size with boundary type. It is clear that the family size effect is only really significant for the double consonants.

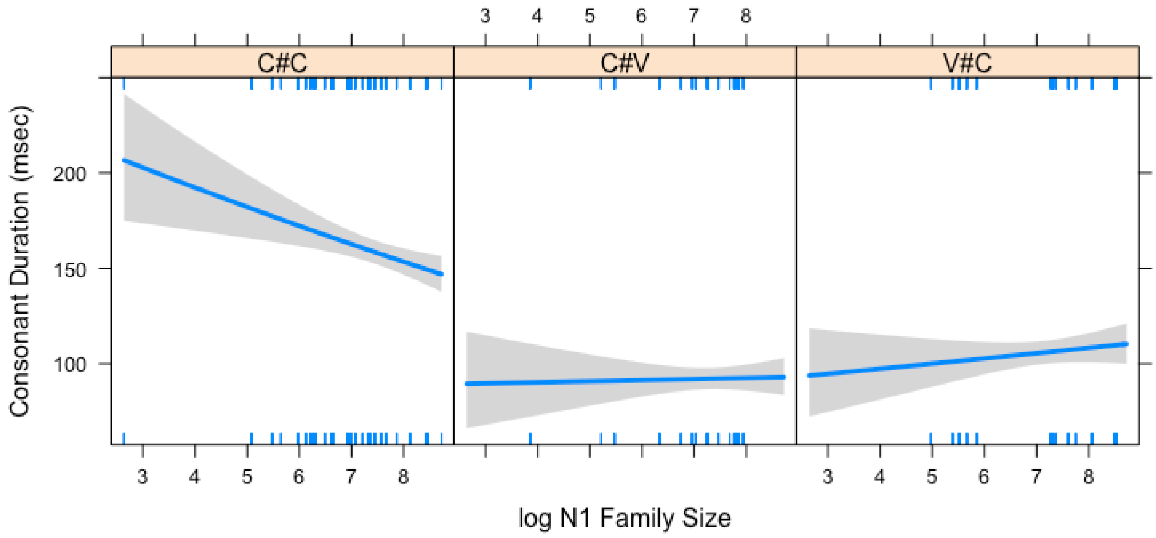

Figure 5: Partial effect of N1 family size by boundary type in final model

The only main effect in our model is a rather weak effect of N2 family size, shown in Figure 6. This goes in the direction predicted by the informativity hypothesis, with larger N2 family sizes going together with shorter consonant durations at the boundary.

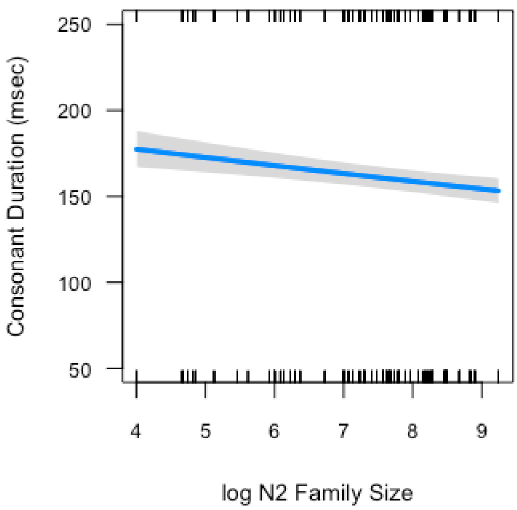

Figure 6: Partial effect of N2 family size in final model

5 Discussion and Conclusion

In this study we have investigated the duration of consonants at the boundary of noun-noun compounds in English. Following up on other studies that have looked at phonetic correlates of morphological structure, the present work tested three hypotheses that have been put forward in other research that has sought to understand morpho-phonetic effects.

In our speech production experiment we tested the duration of three different consonants in three different environments, using a number of variables that gauge effects of segmentability, informativity and paradigm complexity. In addition control variables were used to account for well-known phonetic effects of enhancement or reduction.

The control variables showed the expected effects. For example, higher speech rate led to shorter consonants. Morphological gemination yielded longer consonant durations. This finding is in line with similar results in recent studies (Ben Hedia and Plag, 2017; Ben Hedia, 2018; Plag and Ben Hedia, 2018; Kotzor et al., 2016) on double consonants straddling affix or compound boundaries. With regard to the variables of interest we came up with the following results:

Of the many variables we tested, two were most influential, the family sizes of N1 and N2. Both family sizes showed significant effects on consonant duration. Increasing the family size of N1 leads to shorter durations. This is an effect that can be interpreted as a paradigmatic enhancement effect, since in compounds with a large N1 family size, N2 has little paradigmatic support and is therefore shortened. This effect is analogous to the one observed for Dutch compound interfixes by Kuperman and colleagues (2007). As already mentioned above, the paradigmatic enhancement effect has also been shown to exist in inflectional paradigms. The present study thus provides more evidence that this effect is real.

The N1 family size effect in our study is, however, restricted to geminates. There are several possible reasons for this. It could be related to the fact that there are more tokens of this type in the data, so there is simply more statistical power to enable the effect to reach significance. Or it could be an acoustic effect, with the longer duration of morphological geminates allowing more scope for variation in that duration. Or, as suggested in Section 3.1, it may be that the nature of morphological geminates, straddling as they do a morphological boundary in a single articulatory gesture, makes them especially susceptible to paradigmatic effects.

The N1 family size effect did not go in the direction predicted by the segmentability hypothesis. The other measure used to test that hypothesis, SpellingRatio, did not come out as predicted either. SpellingRatio did not have a significant effect on duration. These two results together suggest that the segmentability of compounds, though predictive of other measures (e.g. of reaction times in priming experiments), is not related to the acoustic duration of the consonant at the compound-internal word boundary. With this outcome the present study joins the list of investigations that failed to replicate the segmentability effects that other studies have found (see Plag and Ben Hedia 2018 for a recent overview).

The N2 family size is best conceptualized as a measure of informativity that taps into the cost of planning the right constituent. Kuperman et al. (2007, p. 2264), for example, state that "[p]lanning upcoming events with a low information load has been shown to predict reduction in the fine phonetic detail of the currently produced elements". In our case the element in question is the consonant at the boundary, and we too find this effect, even if it is rather weak.

It is perhaps surprising that none of the other measures survived in our models. This is partly due to our modeling procedure. To avoid collinearity, we had to exclude one of two or more correlating variables if they had the same effect on the dependent variable. We only kept the most influential one. This does not mean that the other variables might not have an effect on duration too. Entropy, for instance, shows similar effects as the family sizes, only that they are weaker. Across all models, the N1 family size was the most robust effect.

In sum, our study, like some others before, has provided evidence that both informativity and paradigmatic structure influence the fine phonetic details of complex words, in our case the duration of consonants at the compound-internal boundary.

Acknowledgements To be added to final version.

References

Aylett, M. and Turk, A. (2006). Language redundancy predicts syllabic duration and the spectral characteristics of vocalic syllable nuclei. *The Journal of the Acoustical Society of America*, 119:3048–3058.

Baayen, H., Burani, C., and Schreuder, R. (1997). Effects of semantic markedness in the processing of regular nominal singulars and plurals in Italian. In Booij, G. E. and van Marle, J., editors, *Yearbook of Morphology 1996*, pages 13–34. Kluwer Academic Publishers, Dordrecht.

Baayen, R. H. (2010). languageR: Data sets and functions with "Analyzing Linguistic Data: A practical introduction to statistics".

Bates, D., Maechler, M., Bolker, B., and Walker, S. (2014). lme4: Linear mixed-effects models using Eigen and S4.

Bauer, L. (1998). When is a sequence of two nouns a compound in English? *English Language and Linguistics*, 2(1):65–86.

Bell, A., Brenier, J. M., Gregory, M., Girand, C., and Jurafsky, D. (2009). Predictability effects on durations of content and function words in conversational English. *Journal of Memory and Language*, 60(1):92–111.

Bell, M. (2011). At the boundary of morphology and syntax: Noun noun constructions in English. In Galani, A., Hicks, G., and Tsoulas, G., editors, *Morphology and its interfaces*, pages 137–167. John Benjamins, Amsterdam and Philadelphia.

Ben Hedia, S. (2018). *Gemination and degemination in English affixation: Investigating the interplay between morphology, phonology and phonetics*. Dissertation, Heinrich-Heine-Universität Düsseldorf, Düsseldorf.

Ben Hedia, S. and Plag, I. (2017). Gemination and degemination in English prefixation: Phonetic evidence for morphological organization. *Journal of Phonetics*, 62:34–49.

Bermúdez-Otero, R. (2018). Stratal phonology. In Hannahs, S. J. and Bosch, A., editors, *Routledge handbook of phonological theory*, pages 100–134. Routledge, London, UK.

Boersma, P. and Weenink, D. (2014). Praat: doing phonetics by computer.

Bürki, A., Ernestus, M., Gendrot, C., Fougeron, C., and Frauenfelder, U. H. (2011). What Affects the Presence Versus Absence of Schwa and its Duration: A Corpus Analysis of French Connected Speech. *Journal of the Acoustical Society of America*, 130(6):3980–3991.

Cohen, C. (2014a). *Combining structure and usage patterns in morpheme production: Probabilistic eCombining structure and usage patterns in morpheme production: Probabilistic effects of sentence context and inflectional paradigms.* PhD dissertation, University of California, Berkeley.

Cohen, C. (2014b). Probabilistic reduction and probabilistic enhancement. *Morphology*, 24(4):291–323.

Hanique, I. and Ernestus, M. (2012). The Processes Underlying Two Frequent Casual Speech Phenomena in Dutch: A Production Experiment. In *Proceedings of the 13th Annual Conference of the International Speech Communication Association.*

Hay, J. (2007). The phonetics of un-. In Munat, J., editor, *Lexical creativity, texts and contexts*, pages 39–57. Benjamins, Amsterdam / Philadelphia.

Hay, J. B. (2003). *Causes and Consequences of Word Structure.* Routledge, New York and London.

Hay, J. B. and Baayen, H. (2003). Phonotactics, Parsing and Productivity. *Italian Journal of Linguistics*, 1:99–130.

Ji, H., Gagné, C. L., and Spalding, T. L. (2011). Benefits and costs of lexical decomposition and semantic integration during the processing of transparent and opaque English compounds. *Journal of Memory and Language*, 65(4):406–430.

John Coleman, J., Baghai-Ravary, L., Pybus, J., and Grau, S. (2012). *Audio BNC: the audio edition of the Spoken British National Corpus.* http://www.phon.ox.ac.uk/AudioBNC.

Jurafsky, D., Bell, A., Gregory, M., and Raymond, W. D. (2001). Probabilistic Relations Between Words: Evidence from Reduction in Lexical Production. In Bybee, J. and Hopper, P., editors, *Frequency and the emergence of linguistic structure*, Typological Studies in Language, pages 229–254. John Benjamins, Amsterdam.

Kiparsky, P. (1982). From cyclic phonology to lexical phonology. In van der Hulst, H. and Smith, N., editors, *The structure of phonological representations*, pages 131–176. Foris, Dordrecht.

Kotzor, S., Molineaux, B. J., Banks, E., and Lahiri, A. (2016). fake gemination in suffixed words and compounds in english and german. *The Journal of the Acoustical Society of America*, 140(1):356–367.

Kunter, G. and Plag, I. (2016). Morphological embedding and phonetic reduction: The case of triconstituent compounds. *Morphology*, 26(2):201–227.

Kuperman, V., Bertram, R., and Baayen, H. (2009). Processing trade-offs in the reading of Dutch derived words. *Ms.*

Kuperman, V., Bertram, R., and Baayen, R. (2008). Morphological Dynamics in Compound Processing. *Language and Cognitive Processing*, 23:1089–1132.

Kuperman, V., Bertram, R., and Baayen, R. H. (2010). Processing trade-offs in the reading of Dutch derived words. *Journal of Memory and Language*, 62:83–97.

Kuperman, V., Pluymaekers, M., Ernestus, M., and Baayen, H. (2007). Morphological predictability and acoustic salience of interfixes in Dutch compounds. *JASA*, 121:2261–2271.

Kuznetsova, A., Brockhoff, P. B., and Christensen, R. H. B. (2017). lmertest package: tests in linear mixed effects models. *Journal of Statistical Software*, 82(13).

Ladefoged, P. (2003). *Phonetic data analysis.* Blackwell, Malden and Mass.

Lee-Kim, S.-I., Davidson, L., and Hwang, S. (2013). Morphological effects on the darkness of English intervocalic /l/. *Laboratory Phonology*, 4(2):475–511.

Levelt, W. J. M., Roelofs, A., and Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, 22:1–38.

Lindblom, B. (1963). Spectrographic study of vowel reduction. *The Journal of the Acoustical Society of America*, 35(11):1773–1781.

Lõo, K., Järvikivi, J., Tomaschek, F., Tucker, B. V., and Baayen, R. H. (2018). Production of Estonian case-inflected nouns shows whole-word frequency and paradigmatic effects. *Morphology*, 28(1):71–97.

MacGregor, L. J. and Shtyrov, Y. (2013). Multiple routes for compound word processing in the brain: Evidence from EEG. *Brain and Language*, 126(2):217–229.

Milin, P., Filipović Durdević, D., and Moscoso del Prado Martín, Fermin (2009a). The Simultaneous Effects of Inflectional Paradigms and Classes on Lexical Recognition: Evidence from Serbian. *Journal of Memory and Language*, 60(1):50–64.

Milin, P., Filipović Durdević, D., and Moscoso del Prado Martín, Fermin (2009b). The simultaneous effects of inflectional paradigms and classes on lexical recognition: Evidence from Serbian. *Journal of Memory and Language*, 60(1):50–64.

Milin, P., Kuperman, V., Kostić, A., and Baayen, R. H. (2009c). Words and paradigms bit by bit: An information-theoretic approach to the processing of paradigmatic structure in inflection and derivation. In Blevins, J. P. and Blevins, J., editors, *Analogy in grammar: form and acquisition*, pages 214–252. Oxford University Press, Oxford.

Nooteboom, S. G. (1972). *Production and perception of vowel duration: A study of the durational properties of vowels in Dutch.* University of Utrecht, Utrecht.

Plag, I. and Balling, L. W. (2017). Derivational morphology: An integrated perspective. In Pirrelli, V., Dressler, W. U., and Plag, I., editors, *Word knowledge and word usage: A cross-disciplinary guide to the mental lexicon*. de Gruyter Mouton, Berlin.

Plag, I. and Ben Hedia, S. (2018). The phonetics of newly derived words: Testing the effect of morphological segmentability on affix duration. In Arndt-Lappe, S., Braun, A., Moulin, C., and Winter-Froemel, E., editors, *Expanding the Lexicon: Linguistic Innovation, Morphological Productivity, and the Role of Discourse-related Factors*, pages 93–116. de Gruyter Mouton, Berlin, New York.

Plag, I., Kunter, G., Lappe, S., and Braun, M. (2008). The role of semantics, argument structure, and lexicalization in compound stress assignment in English. *Language*, 84(4):760–794.

Pluymaekers, M., Ernestus, M., Baayen, R. H., and Booij, G. (2010). Morphological effects in fine phonetic detail: The case of Dutch -igheid. In Fougeron, C., editor, *Laboratory phonology 10*, Phonology and phonetics, pages 511–531. Mouton de Gruyter, Berlin, New York.

Schreuder, R. and Baayen, H. (1997). How complex simplex words can be. *Journal of Memory and Language*, 37:118–139.

Schuppler, B., van Dommelen, W. A., Koreman, J., and Ernestus, M. (2012). How Linguistic and Probabilistic Properties of a Word Affect the Realization of its Final /t/: Studies at the Phonemic and Sub-Phonemic Level. *Journal of Phonetics*, 40(4):595–607.

Seyfarth, S. (2014). Word informativity influences acoustic duration: Effects of contextual predictability on lexical representation. *Cognition*, 133(1):140–155.

Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27:379–423.

Smith, R., Baker, R., and Hawkins, S. (2012). Phonetic detail that distinguishes prefixed from pseudo-prefixed words. *Journal of Phonetics*, 40(5):689–705.

Sproat, R. and Fujimura, O. (1993). Allophonic variation in English /l/ and its implications for phonetic implementation. *Journal of Phonetics*, 21:291–311.

Tomaschek, F., Plag, I., Baayen, R. H., Baayen, and Ernestus, M. (2018). Modeling the duration of word-Modeling the duration of word-final S in English with Naive Discriminative Learning. *Submitted*.

Torreira, F. and Ernestus, M. (2012). Weakening of intervocalic /s/ in the Nijmegen Corpus of Casual Spanish. *Phonetica*, 69(3):124–148.

Tucker, B. V., Sims, M., and Baayen, R. H. (2018). Opposing forces on acoustic duration. *Submitted for publication*.

Umeda, N. (1977). Consonant duration in American English. *The Journal of the Acoustical Society of America*, 61(3):846.

van Jaarsveld, H. J., Coolen, R., and Schreuder, R. (1994). The role of analogy in the interpretation of novel compounds. *Journal of Psycholinguistic Research*, 23:111–137.

van Son, R. J. and Pols, L. C. (2003). Information Structure and Efficiency in Speech Production. In *Proceedings of Eurospeech-2003*, pages 769–772, Geneva and Switzerland.

van Son, R. J. and van Santen, J. P. (2005). Duration and spectral balance of intervocalic consonants:A case for efficient communication. *Speech Communication*, 47:100–123.

Wedel, A., Kaplan, A., and Jackson, S. (2013). High functional load inhibits phonological contrast loss: A corpus study. *Cognition*, 128(2):179–186.

Zimmermann, J. (2016). Morphological status and acoustic realization: Is there a difference between Bra[d] Pitt and a grille[d] cheese omelet, or between Kate Mo[s] and killer robot[s]? In Wahle, J., Köllner, M., Baayen, H. R., Jäger, G., and Baayen-Oudshoorn, T., editors, *Proceedings of 6th Conference on Quantitative Investigations in Theoretical Linguistics (QITL-6)*, Tübingen. Universität Tübingen.

Zwitserlood, P. (2007). The role of semantic transparency in the processing and representation of Dutch compounds. *Language and Cognitive Processes*, 9(3):341–368.