Interaction in Phonological Variation:

Grammatical Insights from a Corpus-Based Approach

Jeffrey Lamontagne

Linguistics Department

McGill University, Montréal

September 28th, 2020

A thesis submitted to McGill University in partial fulfillment for the requirements of the degree of Doctor of Philosophy

Abstract

This thesis probes the role of grammatical structure and cognitive processing in phonological variation using a corpus-based approach. Study 1 investigates variable prominence assignment through a corpus study of Laurentian French. Prominence is conventionally described as being assigned to phrase-final syllables in French, but previous work has shown that this is not always the case. Using mixed-effects linear regression to analyse corpus data from Laurentian French, I test the hypothesis that prominence assignment signals weight contrasts. The results support this hypothesis, with both codas and long vowels attracting prominence away from final syllables, particularly when the final syllable is open. In terms of phonetic realisation, pitch and amplitude are the primary cues to prominence, consistent with descriptions in the literature, with amplitude additionally manipulated to signal weight but not phrasal prominence. These findings suggest that prominence is best formally expressed as a pitch accent due to its attraction to word-level properties. I propose that phonetic and phonological evidence converge on this analysis of the prominence system and can inform our interpretation of results from previous perceptual studies.

Study 2 examines the tense/lax quality of high vowels in non-final syllables in Laurentian French. High-vowel tenseness has evaded large-scale analysis in production due to interactions among multiple optional processes, as well as challenges posed by acoustic similarity between tense and lax high vowels. In final syllables, high-vowel tenseness is fully predictable. In non-final syllables, however, high-vowel tenseness has been proposed to be determined by several phonological processes triggered by syllable structure, morphological structure, and featural constraints. In this study, a forced aligner is trained to classify high-vowel tenseness using tokens of final-syllable high vowels, where tenseness is categorically predictable. The forced aligner is then used to classify the tenseness of high vowels in non-final syllables. The results of mixed-effects logistic regression provide support for only some processes. I conclude that high-vowel tenseness is determined by phonological processes rather than phonetic pressures, that tenseness is sensitive to morphological structure and, contrary to previous proposals, that learners receive sufficient input to generate a grammar of high-vowel tenseness despite variability and opacity.

Study 3 investigates how cognitive processing motivates phonological variability. It tests the hypothesis that the variable capacity for speakers to plan the context of cross-word phonological processes results in variability (Production Planning Hypothesis; PPH). Implicit to the PPH is the prediction that processes where the trigger follows the target will be more sensitive to speakers' ability to plan phonological contexts than processes where the trigger precedes the target because the trigger is only consistently expected to be planned early in the latter case. I employ a forced aligner to automate detection of deletion in /a#e/ and /e#a/ sequences in conversational Madrid Spanish and use Bayesian mixed-effects multinomial regression to test the role of factors associated with ease of production planning (lexical frequency, conditional probability). I show that deletion where the trigger precedes the target shows significantly smaller and fewer production planning effects, consistent with the PPH.

The three corpus studies that comprise this thesis shed light on how grammatical variability is associated with phonological structure and cognitive processing. They particularly show how techniques for corpus analysis uncover systemic patterns in what previously appeared to be unstructured variation. I conclude with discussion of insights provided by the three studies into the relationships between phonological variation and both prosody and morphology.

Abrégé

Cette thèse utilise des données de corpus pour explorer le rôle de la structure grammaticale et du traitement cognitif dans la variation phonologique. L'étude 1 examine l'attribution variable de la proéminence à travers une étude du français laurentien. La proéminence est classiquement décrite comme étant attribuée à la dernière syllabe des phrases en français, mais des travaux récents ont démontré que ce n'est sans exception. En utilisant une régression linéaire à effets mixtes, je teste l'hypothèse selon laquelle la proéminence signale le poids phonologique. Les résultats soutiennent cette hypothèse : les codas et les voyelles longues attirent la proéminence, surtout quand la syllabe finale est ouverte. En termes de réalisation phonétique, la hauteur et l'amplitude sont les principaux indices de la proéminence, avec l'amplitude manipulée pour signaler le poids mais pas la proéminence phrasale. Ces résultats suggèrent que la proéminence est formellement un accent de hauteur attiré par des propriétés lexicales. Je propose que les preuves phonétiques et phonologiques convergent vers cette analyse du système de proéminence et puissent informer notre interprétation des résultats d'études perceptuelles antérieures.

L'étude 2 examine le relâchement des voyelles hautes dans les syllabes non finales en français laurentien. Ce phénomène a échappé à l'analyse à grande échelle dans la production en raison d'interactions entre processus optionnels, ainsi que des défis posés par la similitude acoustique entre allophones. En syllabe finale, le relâchement des voyelles hautes est entièrement prévisible. Cependant, en syllabe non finale, plusieurs processus phonologiques ont été proposés pour prédire le relâchement de ces voyelles – déclenchés par la structure syllabique, par la structure morphologique et par des contraintes sur les traits. Dans cette étude, un aligneur automatique est entraîné à classer les voyelles hautes en syllabe finale. L'aligneur classe ensuite le relâchement des voyelles hautes en syllabe non finale. Une analyse de régression logistique à effets mixtes ne soutient que certains processus. J'en conclus que le relâchement des voyelles hautes est proprement phonologique plutôt que phonétique, que la tension est sensible à la structure morphologique et, contrairement à une proposition précédente, que l'input que

reçoivent les apprenants est suffisante pour générer une grammaire du relâchement des voyelles hautes malgré la variabilité et l'opacité.

L'étude 3 examine comment le traitement cognitif engendre la variabilité phonologique. Je teste l'hypothèse que la capacité variable à planifier le contexte des processus phonologiques entraîne la variabilité (l'hypothèse de planification langagière; HPL). L'HPL prédit implicitement que les processus où le déclencheur suit la cible seront plus sensibles à la la facilité à planifier le contexte phonologique que les processus où le déclencheur précède la cible parce que le déclencheur ne sera planifié tôt de façon fiable que dans ce dernier cas. J'utilise un aligneur automatique pour localiser l'amuïssement vocalique dans les séquences /a#e/ et /e#a/ en espagnol de Madrid et j'utilise la régression multinomiale bayésienne à effets mixtes pour tester le rôle des facteurs associés à la facilité de planification de la production (fréquence lexicale, probabilité conditionnelle). Je montre que l'amuïssement où le déclencheur précède la cible montre des effets de planification de la production nettement plus faibles, conformément à l'HPL.

Les trois études de cette thèse mettent en lumière la manière dont la variabilité grammaticale est associée à la structure phonologique et au traitement cognitif. Ils montrent en particulier comment les techniques d'analyse de corpus témoignent de tendances systémiques dans ce qui semblait auparavant être de la variation non structurée. Je conclus en discutant des relations entre (a) la variation phonologique et (b) la prosodie et la morphologie.

Table of Contents

List of Figures

List of Tables

Acknowledgements

I cannot thank both of my supervisors enough for their guidance, support and encouragement. Hea*the*r Goad was a gift from day 1 in *the* program. She offered incredible support for research and showed unparalleled excitement when teaching us at what sometimes felt like *the* wee hours of *the* morning. Fur*the*rmore, and perhaps even more importantly, she reinforced that what happens in non-academic life deserves attention and time. Her careful and thoughtful commenting was always appreciated, and through *the*m I discovered all new idiosyncrasies and grammatical nuggets to keep in mind. Meetings "may" often have been at risk of running over, but *the* results of those meetings were always improvements, whe*the*r that be to a paper, to a course, or to my frame of mind. Francisco Torreira joined the department later into my program, but nonetheless contributed considerably not only to academic interests, but also to how I view our field. Shana Poplack had previously instilled a little voice in my head asking "*Who cares?*" when contemplating results and generalisations, but Francisco highlighted that the audience for that question doesn't always need to be academic. It is common for linguists to get wrapped up in theoretical implications and in overarching questions, but we should not forget to provide descriptions of language use that – ideally – may be of interest to speech communities.

I would additionally like to thank the McGill Linguistics department, particularly Morgan Sonderegger who helped supervise my first research paper and who guided me through so much in R, in Python, and in stats more broadly. Thank you also to P* Reading Group members throughout these years for lively and thought-provoking discussion. Thank you to Meghan Clayards and to Michael Wagner for discussion at so many points in and out of class. Finally, thank you to Andria De Luca and to Giuliana Panetta, who always seem to have their hand in things running smoothly in the department.

Thank you to Marie-Hélène Côté, Jeff Mielke, and Andrés Salanova, whose support and encouragement at the University of Ottawa played a pivotal role in shaping my love for linguistics and for research. Their excitement and teaching set me on this path, and I couldn't imagine where I would be without them. Thank you to Peter Milne, whom I met at the

University of Ottawa and then got to know better as an ever-supportive colleague while at McGill.

Thank you to my fellow students at McGill – notably Bing'er Jiang, Colin Brown, Dejan Milačić, Donghyun Kim, Emily Kellison-Lynn, Francesco Gentile, Gui Garcia, Henrison Hsieh, Hye-Young Bang, James Tanner, Mathieu Paillé, Oriana Kilbourn-Ceron, Sepideh Mortazavinia, and Yeong Woo Park. At various points, we've shared offices (officially or not) or time. The discussion and commiseration and gossip made even the most frustrating days more enjoyable. For these same things, thank you to Gretchen McCulloch, Hadas Kotek, Meaghan Fowlie and Nico Baier. It was a pleasure to relax – or to redirect stress! – with all of you.

And, of course, I am thankful to my family. These years have been both challenging and wonderful in so many ways, and I couldn't have gotten through them without you. (And Steven, enjoy that they've now been trained not to ask the taboo questions… well, as often, at least.) Thank you also to Raf, Rémi, Rose, and Sammy, who have offered joy and distraction during these years.

Finally, I would like to acknowledge the financial support of SSHRC Joseph Armand Bombardier Doctoral Fellowship (752-2014-1969).

Contribution of Authors

The studies that comprise this thesis were prepared as manuscripts for publication elsewhere. I am the primary author of each manuscript.

Chapter 2 is being prepared for publication as an article co-authored with Prof. Heather Goad. I was responsible for the original design and conception of the study, as well as the data preparation and analysis. The statistical analysis was performed by me in consultation with Prof. Morgan Sonderegger, and the manuscript was prepared in collaboration with Prof. Goad.

Chapter 3 is being prepared for publication as an article co-authored with Dr. Peter Milne. Dr. Milne and I designed the study together, and we prepared the data together. Dr. Milne made the adaptations to the forced aligner, and I performed the data analysis and prepared the manuscript.

Chapter 4 is being prepared for publication as an article co-authored with Prof. Francisco Torreira. I designed the study, prepared the data, and performed the data analysis. The manuscript was prepared in collaboration with Prof. Torreira.

Chapter 1

Introduction

Variation poses considerable challenges to phonological theory because of the difficulty accounting for a single underlying form yielding multiple possible surface forms. This is further complicated by the lack of invariance problem in phonetics (e.g. Blumstein and Stevens, 1981; Appelbaum, 1996) that gradient phonetic realisations are inherently variable. Variability in phonetic and phonological realisations is often characterised as free variation because a speaker may choose any pronunciation without changing the meaning of a word or sentence (e.g. Coetzee, 2009; Fruehwald, forthcoming). However, so-called free variation is not entirely free; words may pattern differently as a result of lexical properties (e.g. Gimson, 1969) or prosodic context (e.g. Fougeron, 1999), and speakers may modulate variation based on their social context and affiliations (e.g. Fischer, 1958; Labov, 1972). Adopting a term from Weinreich, Labov and Herzog (1968), the presence of "orderly heterogeneity" (structured variation) may even be inevitable under the assumption that variation largely emerges from phonetic differences in articulation and perception (see Baker et al., 2011; Beddor et al., 1986) and that all change passes through a stage of variation (Weinreich et al., 1968). The goal of this thesis, then, is to probe variation to determine what it reveals about the grammar of a language and the mechanisms underlying speech production and perception.

Recent work has especially benefited from the creation of large-scale speech corpora and the development of techniques for analysing those corpora. Corpus analysis techniques allow researchers to elucidate relatively subtle structure in variation, drawing on larger amounts of data than previously possible. Amongst the most influential advancements in corpus analysis is the development of forced alignment, which is a technique to automate the segmentation and annotation of audio files based on preexisting transcriptions. Data preparation techniques coupled with advancements in statistical analysis techniques (e.g. mixed models to account for

idiosyncratic differences between speakers or words) to better understand individual or group behaviours. As I illustrate in this thesis, these tools can be leveraged to serve additional purposes. For example, forced alignment provides a method of classifying pronunciations into allophones without manual input. Additionally, statistical models can illustrate successful grammar generation by a learner (for an example and discussion, see Goldwater and Johnson, 2003; Hayes and Wilson, 2008; Wilson, 2006).

Existing work confirms that structured variation can provide insight into both the grammar and the mechanisms involved in speech production. The examination of acoustic cue usage for stop contrasts informs our understanding of the phonological contrasts that are conveyed through those acoustic cues (Bang, 2017; Chodroff, 2017). Similarly, the deletion rate of English coronal stops differs based on the morphological status of the stop and therefore can be levied to determine whether a speaker has begun to parse irregular verbs as morphologically complex (e.g. *kept* as *keep* + *-d*; Guy and Boyd, 1990). Coronal stop deletion in English has further been studied to demonstrate that the ease of phonological encoding modulates the application of phonological processes (Kilbourn-Ceron, 2017; Kilbourn-Ceron et al., 2017; Tanner, 2017). In short, variation can inform our understanding of grammar and how speakers process speech.

This thesis applies a corpus-based approach to the examination of three cases of variation to show how variation emerges from phonological structure and cognitive processing. Through studying this variation, we can therefore gain a better understanding of aspects of speech that may not be evident from introspective judgments. The three studies are introduced in section 1.1.

1.1 Case Studies

I approach the contributions that an examination of structured variation makes to the study of language through three case studies. In each case, the presence of variation is well attested. It is generally unclear, however, what motivates variation in the phenomena under focus and, in one case (Study 1), the existence of variation remains controversial. Study 1 (§1.1.1) focuses on prominence assignment in Laurentian French, investigating motivations for main prominence to variably surface on either the penultimate or final syllable of phrases. Study 2 (§1.1.2) turns to the tense/lax quality of high vowels in Laurentian French, examining the complex set of

phonological processes that the literature has proposed in order to account for high vowels surfacing as tense or lax. Finally, Study 3 (§1.1.3) probes the resolution of cross-word vowel hiatus in Madrid Spanish, testing cognitive explanations for variability in its application.

1.1.1 Prominence in Laurentian French

In French, main prominence is conventionally described as being assigned to the final syllable of phrases (e.g., Delattre, 1939; Grammont, 1914; Jun and Fougeron, 1995, 2000, 2002; Martin, 1987; Pasdeloup, 1990; Selkirk, 1972). It is usually agreed that this prominence is primarily reflected through the presence of a pitch peak, with longer durations as a secondary cue (e.g. Jun and Fougeron, 1995, 2000). Some varieties of French are described as also making use of amplitude (Alsace, Burgundy, Jura, Lorrace, Normandy; Carton et al., 1983) and speakers described as speaking "Standard French" have been found to readily use amplitude in perceiving prominence (Li et al., 2017; Schwab and Llisteri, 2012). While cue usage may be subject to variation, the more important point of variation for Study 1 is the syllable to which prominence is assigned. Across the French-speaking world, prominence has been observed to *shift* to the penult even if we restrict the scope to studies that have examined the use of pitch in native speakers' productions: in Parisian French and other northern Hexagonal varieties (Carton et al., 1983; Goldman and Simon, 2007; Simon, 2011), Belgian French (Bardiaux and Mertens, 2014; Goldman and Simon, 2007; Simon, 2004, 2011); Laurentian French (Thibault and Ouellet, 1996), and Swiss French (Avanzi et al., 2011; Goldman and Simon, 2007).

Variation in prominence assignment has been attributed to multiple factors, generally without comparing speakers or speech contexts to test hypotheses. Post (2000) proposed that specific contours are targeted in shifting the location of prominence and Avanzi et al. (2011a) looked to speech rate differences to explain variation. The most common explanation, however, is contact with a dialectal substrate (Carton 1979) or with other languages (e.g. Kaminskaïa and Poiré, 2012; Mamode, 2015; Sichel-Bazin, 2012; Sichel-Bazin et al., 2011). Relating to the possibility of dialectal differences, Martin (2004) suggests that prominence shift reflects archaic tendencies. There is reason to expect that prominence shift might be a conservative phenomenon: Avanzi et al. (2011b) observe that older speakers produce prominence shift more often than younger

speakers. Furthermore, a comparison of studies examining European speakers born since around 1900 suggests that prominence shift is rarifying in Northern France (based on duration measures due to data limitations; Boula de Mareüil, 2008; Martin, 2011), though this is not necessarily the case across sociolects (Fagyal, 2003). The question remains whether and how variation in prominence assignment is structured: *why* do speakers shift prominence only in a subset of cases given that they shift at all?

Study 1 uses corpus data to test the hypothesis that (Laurentian) French speakers exhibit syllable weight contrasts and that prominence shift results from variable weight sensitivity. It has been observed across languages that heavy syllables attract stress (Gordon, 2014; Prince, 1990). Weight sensitivity varies in some ways, however, in that not all phonological systems treat the same types of syllables as heavy (Hayes, 1995). The most common source of heavy syllables is long vowels. The second possible source is codas, whereby a closed syllable will be treated as heavy and an open one will not be. Regarding French weight, the language can be analysed as having both long vowels and weight-bearing codas. Laurentian French is examined in this case study because it has a particularly large and stable vowel inventory relative to other varieties of French (cf. Côté, 2012). Study 1 therefore tests whether variable prominence assignment in French results from variable weight sensitivity, thereby additionally testing whether French vowels can be considered to differ in weight.

1.1.2 Tense/Lax Quality in Laurentian French

The second study turns to another phenomenon in Laurentian French that has challenged analysts because of rampant variability and complexity. In addition to having a relatively large inventory of vowel contrasts, Laurentian French is characterised by phonological distinctions in the tense/lax quality of high vowels. In final syllables, the tense/lax quality is fully predictable: high vowels are tense in open syllables (and syllables closed by a handful of consonants) and lax elsewhere. In non-final syllables, however, the presence of variable high-vowel laxing has been described since Dumas (1983, 1987). Testing this variability has posed a challenge: while corpora of native Laurentian French speakers are available (e.g. Canavan and Zipperlen, 1996; Poplack, 1989; Sankoff et al., 1976; Thibault and Vincent, 1990; Vincent et al., 1995), the

acoustic similarity between tense and lax high vowels has rendered analyses of production infeasible (e.g. Arnaud et al., 2011; Sigouin, 2013). As such, the literature has largely relied on perceptual results from Poliquin (2006) in formal accounts of Laurentian French high vowels in non-final syllables (e.g. Bosworth, 2011; Fast, 2008; Poliquin, 2006). The resulting analyses have been complex; several optional phonological processes have been proposed, most notably laxing in closed syllables, laxing disharmony, tensing due to coda resyllabification, and multiple types of laxing harmony. The situation is further complicated by these processes interacting with each other and applying at different stages of the derivation to result in variable opacity. Arising from the complexity of these interactions, the proposed poverty of the stimulus has been suggested to render it impossible for learners to converge on a single grammar of tense/lax quality in Laurentian French (Poliquin, 2006).

Study 2 leverages developments in corpus tools to classify the tense/lax quality in non-final syllables, thereby enabling the examination of production data to test phonological processes proposed in the literature. A forced aligner was trained to distinguish tense and lax high vowels based on final syllables (where allophones are fully predictable), then that forced aligner was run to classify high vowels in non-final syllables (where most phonological processes are optional) to elucidate the grammar of tense/lax quality. In this study, I test two aspects of the grammar of the tense/lax quality in Laurentian French high vowels. First, I investigate whether the difficulty accounting for the tense/lax quality in high vowels is due to misattributing phonetic pressures (coarticulation, reduction) as being phonological processes. Second, I identify the grammar of the tense/lax quality that emerges from community-level data (i.e., the grammar a learner is expected to generate). Taken together, examination of the tense/lax quality provides multiple potential insights into the structure of the grammar.

1.1.3 Cross-Word Hiatus in Madrid Spanish

The third study turns to other explanations for variation, notably cognitive processing. In Spanish, vowel-vowel sequences are avoided both within and across words (e.g. Aguilar, 1999; Baković, 2006; Hualde and Chitoran, 2003; Souza, 2010; Vuskovich, 2006). One way that Spanish speakers resolve hiatus is through deletion, which can target either of the two vowels.

When /a/ and /e/ are adjacent in /e#a/ and /a#e/ sequences, there is a strong preference to delete /e/ and preserve /a/ (e.g. Garrido, 2008), suggesting that /a/ serves as the trigger for /e/ deletion when the two vowels are adjacent. This potential for V1 to trigger deletion of V2 or for V2 to trigger deletion of V1 provides a rare opportunity to test for asymmetries in directionality in what appears to be a single phonological process. Study 3 therefore investigates potential planning effects that are evidenced by directional asymmetries in cross-word processes.

The Production Planning Hypothesis (PPH; Wagner, 2012) states that the ability to phonologically encode words mediates the application of cross-word phonological processes, which results in variability. Under this hypothesis, processes will be variable when the trigger for a process may not be planned in time for the process to apply to the target. The window for phonological encoding may be surprisingly limited (Levelt et al., 1999; Wheeldon, 2012) and therefore two adjacent words may not both be planned together. Factors that are associated with the ease of production planning (e.g. syllable count, lexical frequency, conditional probability, speech rate, and strength of an intervening prosodic boundary) are therefore expected to correlate with the rate of phonological process application.

Thus far, the PPH has predominantly been tested in cases where the trigger of a phonological process followed the target – i.e., cases where planning is expected to be limited. In Tokyo Japanese, the devoicing of a high vowel before a voiceless consonant is sensitive to both speech rate and lexical frequency (Kilbourn-Ceron, 2017). Similarly, English coronal stops are more likely to flap word-finally when the following vowel-initial word is more frequent or when the stop-final word is frequent, while stops glottalise domain-finally regardless of the shape of the upcoming word and therefore this process is not affected by the lexical frequency of the upcoming word (Kilbourn-Ceron, 2017; Kilbourn-Ceron, Wagner, Clayards, 2016). Finally, coronal stops delete pre-consonantally in English, but this deletion is less likely when a stronger prosodic boundary intervenes (Tamminga, 2018; Tanner et al., 2017).

Study 3 tests production planning effects in cross-word hiatus resolution in Madrid Spanish. Once again, a forced aligner is employed for automated classification, this time to code /a#e/ and /e#a/ sequences as having undergone V1 deletion, V2 deletion, or no deletion. Following the PPH, it is expected that /e/ deletion will be more strongly correlated with production planning factors (lexical frequency, conditional probability) in /e#a/ sequences, where planning the trigger

/a/ is expected to be variable, compared to in /a#e/ sequences, where the trigger /a/ is expected to be planned before the target /e/. An asymmetry of this sort would provide strong support for the PPH because alternative explanations (e.g. the reduction of high-frequency words) do not predict both that the frequency of *non*-target words affects process application and that frequency effects are modulated by the order of the trigger and the target.

1.2 Summary

This thesis employs a corpus-based approach to investigate how structured variation reveals properties of cognitive processing and the structure of the grammar. Taken together, the three studies identify motivations for phonological variation. Furthermore, they demonstrate that the study of variation sheds light on the shape of linguistic systems. Variation is thus not only an inevitable fact about language, but also a tool for its study.

Chapter 2

Prominence in Laurentian French

2.1 Introduction

Final syllables are conventionally characterised as bearing the main prominence in French (e.g., Delattre, 1939; Grammont, 1914; Jun and Fougeron, 1995, 2000, 2002; Martin, 1987; Pasdeloup, 1990; Selkirk, 1972). This prominence is reflected through longer durations and higher pitch peaks than are found on adjacent non-prominent syllables (e.g., Jun and Fougeron, 1995, 2000). For example, the final syllable in *patronne* /patʁɔn/ 'boss (FEM)' is the longest and has the highest pitch of the two syllables, which we indicate with the diacritic for stress: [pa'tʁɔn].[1] In contrast to languages like English with lexical stress, prominence in French is not a word-level phenomenon, but is instead assigned at the phrasal level. In a phrase with two lexical words like *la future patronne* 'the future boss (FEM)', only the final syllable is prominent in French: [la fytyʁ pa'tʁɔn].[2] In the corresponding phrase in English, however, both the adjective *future* and the noun *boss* are stressed: [ðə ˌfju:ʧəˇ'bɑs].

The observation that the domains of prominence in French and English are different suggests that prominence serves different functions in the two languages. In French, this cue allows interlocutors to easily and reliably recover the right edge of phrases and therefore reduce the risk of ambiguity (e.g., Mertens, 2006; Vaissière, 2010), for example distinguishing between

[1] Our transcriptions reflect pronunciations in Laurentian French, the dialect under examination. For example, the French rhotic is transcribed as [ʁ] because this is the acoustic realisation that is most common in this variety (Côté and Saint-Amant Lamy, 2012).

[2] While French also has an optional phrase-initial rise (e.g., with low pitch on [la] and high pitch on [fy] in *la future patronne*), we will neither discuss nor transcribe this secondary prominence in the current analysis because we focus only on the right edge. The final high tone is the tone target that is preserved in cases where the phrase has too few syllables to realise both the initial and the final rises (Jun and Fougeron, 2002), meaning that the final high tone can successfully be isolated for study in this paper. Henceforth, we use the term *prominence* to exclusively refer to prominence assigned from the right edge. In addition, the prominence realised on the penult in cases of shift has been shown to be distinct from emphatic stress in Laurentian French (Thibault and Ouellet, 1996), whereas this is not known for antepenultimate prominence.

adjectival modifiers and reduced relative clauses like in the one-phrase parse *la patronne responsable* [la patʁɔn ʁɛspɔ̃'sab(l)], in which the boss is a responsible person, and the two-phrase parse *la patronne responsable* [la pa'tʁɔn ʁɛspɔ̃'sab(l)], in which the boss is responsible for something in particular. However, one challenge for the view that prominence only serves to mark phrasal domains in French is that the cues to prominence do not strictly fall on the final syllable; instead, they often fall on the penult even when the final syllable does not contain a schwa (invisible to prominence assignment; e.g., Garde, 1968, Prieto et al., 2005). This has been observed across varieties of the language: Parisian French and other northern varieties (Carton et al., 1983; Goldman and Simon, 2007; Simon, 2011), Laurentian French (Thibault and Ouellet, 1996), Swiss French (Avanzi et al., 2011; Goldman and Simon, 2007), and Belgian French (Bardiaux and Mertens, 2014; Goldman and Simon, 2007; Simon, 2004, 2011). For example, in *le garçon* 'the boy', the penult can be realised with longer duration and higher pitch than the final syllable: [lə 'gaʁsɔ̃]. This *shift* in the location of prominence suggests that prominence is sensitive to considerations other than just phrasal domain edges. The goal of this paper is to investigate the factors that condition prominence shifts of this sort.

A plausible motivation for prominence shifts is that speakers are enhancing word-level properties instead of – or in addition to – phrase edges. Weight (i.e., the contrast between light and heavy syllables), whether the result of a coda or a long vowel, attracts prominence across languages (Prince, 1990) and could therefore be a relevant word-level property for French because the relative weight of syllables is computed in the domain of the word. If weight plays a role in predicting the location of prominence in French, then the penultimate syllable in *garçon* should be more likely to be prominent than the penultimate syllable in *patronne* because [gaʁ] is closed (and therefore potentially heavy) while [pa] is not. Thus far, the evidence suggesting that weight is responsible for prominence shifts is limited to a study of naive listener judgments (Paradis and Deshaies, 1990), a study looking at the effects of long vowels on pitch contours (Thibault and Ouellet, 1996), and evidence from phonological patterns (Armstrong, 1999; Scullen, 1997).

In this paper, we test the hypothesis that prominence assignment is sensitive to differences in relative weight, consistent with how weight interacts with prominence cross-linguistically. We use mixed effects linear regression to test the effects of prosodic phrasing, of vowel weight and

of coda weight on the realisation of prominence in read speech collected from the Saguenay (Quebec) survey in the Laurentian sub-corpus (Côté, 2014, 2015) of the *Phonologie du français contemporain* (PFC) corpus (Durand et al., 2002, 2009; http://www.projet-pfc.net/). Laurentian French (also referred to as Canadian French, Quebec French, and Québécois) was chosen as it presents more heavy syllables than many other varieties due to having conserved a large number of vowel contrasts from an earlier stage of the language. Furthermore, vowel length is regularly enhanced through diphthongisation (Côté, 2012; Dumas, 1974), facilitating detection of the long-short contrast. In sum, this dialect is a good test case to systematically examine and better understand the relationship between weight and the prominence shift sporadically observed in the literature. We will show that weight effects are associated with the same cues as those manipulated to mark phrase boundaries and that these weight effects provide a motivation for the variable shift in prominence observed. Based on these results, we will argue for a revised interpretation of prominence in French: we will introduce lexical prominence (marked using amplitude and duration), which variably attracts phrasal prominence (marked using pitch, as in the conventional description).

2.2 Theoretical Context

We situate the current study within the context of previous work on prominence. The first topic, discussed in section 2.2.1, is cross-linguistic evidence for weight contrasts and their interaction with prominence. In this section, we additionally discuss weight in French, in order to identify the contexts in which prominence shifts are likely to occur. In section 2.2.2, we turn to the acoustic cues that have been shown to signal prominence, as discussed in the literature on French, in order to guide our examination of the PFC corpus data. We then proceed to a discussion of prosodic domains in section 2.2.3. This is crucial for the present analysis as it identifies the particular domains whose right edge is typically assumed to be marked through prominence, and departures from the baselines established for prosodic domains could indicate an interaction with weight.

2.2.1 Weight

Cross-linguistic comparisons have shown that weight and prominence often interact; in languages with lexical stress, for example, heavy syllables attract stress (Gordon, 2014; Prince, 1990). Using stress assignment in English verbs as an example (Halle, 1973; Hayes, 1982; Liberman and Prince, 1977), once we take into account that final consonants are extrametrical (they do not affect stress assignment), we observe that the final syllable is assigned stress if it contains a complex coda (e.g., *exíst*) or a long vowel (e.g., *eváde*); otherwise the penult is stressed (e.g., *cóvet, hárden*). Example words presented in the literature on French to illustrate prominence shift typically show a similar pattern of weight-sensitivity: prominence is word final unless the penult is heavy and the final syllable is light. We draw upon examples from the literature on French as well as cross-linguistic observations about the interaction of weight and stress to form the hypothesis in (1):

(1) Hypothesis:
French prominence assignment is sensitive to weight.

Not all phonological systems treat the same types of syllables as heavy (Hayes, 1995). The most common source of heavy syllables across languages is vowels: long vowels are heavy, and therefore can attract prominence away from syllables with short vowels. The second possible source of weight is codas, with closed syllables patterning as heavy and therefore being able to attract prominence away from light open syllables.

If heavy syllables attract prominence in French, the question arises as to what counts as heavy: are codas weight-bearing or not, and are there phonemically long vowels in the language? We begin by discussing the status of codas, a term which we use to cover both word-medial rhymal dependents and word-final consonants. French has codas in both positions, as we can see in *marquis* [maʁ.ki] ‘marquis’ and *canal* [ka.nal] ‘canal’, so a related question concerns whether medial and final codas both pattern as heavy. Some authors have analysed word-final codas in French as the onsets of syllables with empty nuclei (Dell, 1995); consistent with this, consonants in this position have an onset profile and clusters with rising sonority are observed word-finally, as in *mettre* [mɛtʁ] ‘to put’, paralleling what are indisputably branching onsets in non-final position in the language (Dell, 1995).

Since onsets do not typically contribute weight[3] and empty nuclei are by definition weightless, final codas are not expected to attract prominence under this analysis, independent of the status of medial codas. In contrast to this view, we analyse final consonants as true codas (i.e., as dependents in the rhyme), based on two observations. One, vowels in final syllables are affected by following consonants in ways that are expected if these consonants are in coda position: for example, final codas are strongly associated with laxing of mid vowels, and in Laurentian French high vowels lax in closed syllables as well (Poliquin, 2007).[4] Two, impressionistic observations regarding the profile of words that have been given as examples of prominence shift in the literature lead us to expect that speakers will be less likely to shift prominence off of the final syllable in a word like *canal* that has a final coda compared to one like *marquis* where the penult is closed but the final syllable is open, which suggests that the final consonant in the former is a weight-bearing coda.

Turning to vowels, French has a relatively large vowel inventory, which is generally described as including both light (short) and heavy (long) vowels even though these contrasts are predominantly realised through quality differences in contemporary French (e.g., Walker, 1984).[5] For example, the upper mid vowel /o/[6] is described as being heavy, which can be seen in final closed syllables in a word like *côte* [ko:t] 'hill' where it is realised as long; in Laurentian French this length can be reinforced by diphthongisation, as in [ko͡ʊt] (e.g., Côté, 2012; Dumas, 1974). This pattern for /o/ can be contrasted with the pattern found for the lower mid vowel /ɔ/, like in the word *cote* [kɔt] 'code', which is not realised as long or diphthongised and is therefore analysed as light. We adopt the position that vowel quality differences reflect length differences, and thus, we expect that prominence will shift inwards more often in a word like *côté* /kote/

[3] Recent work (e.g., Gordon, 2005; Ryan, 2016) has shown that in some languages onsets can contribute weight as well but these effects are small. We leave the possibility that onsets contribute to syllable weight in French for future work.

[4] For Laurentian French, closed-syllable laxing does not productively affect the mid vowels, but the lexicon shows a robust pattern for mid vowels in closed final syllables to be lax, while mid vowels in open final syllables will tend to be tense (Lamontagne, 2014).

[5] We generally use *light* and *heavy* to refer to the classification of vowels rather than *short* and *long* because the phonetic realisations of vowels in Laurentian French predominantly involve spectral differences and not durational differences (as mentioned in the text), and because we want to distinguish between the phonetic measurements of duration used in this paper and the phonological category of weight.

[6] Following typical conventions for French, we transcribe the long vowels – the upper mid vowels /e ø o/, the low back vowel /ɑ/, and the nasal vowels /ɛ̃ œ̃ ɔ̃ ɑ̃/ – without a length diacritic since, as mentioned in the text, the primary cues to the weight contrast are vowel quality and not duration.

'side' than in a word like *coté* /kɔte/ 'coded (as)' because it is attracted to the heavy penult vowel in *côté*, but not to the light penult vowel in *coté*.

This study treats weight from any given source as binary; vowels are light or heavy; open syllables are light (when containing a short vowel) and closed syllables are heavy. While recent work shows that weight in some Romance languages is, instead, gradient (Garcia, 2017), we consider binarity a necessary simplification for our statistical models given the number of tokens being examined, leaving the question of whether gradient weight can be motivated for French to future work. One consequence of this is that syllables with both a long vowel and a coda, for example *côte* [koːt]~[koʊ̯̀t] 'hill', will normally be considered alongside syllables that are heavy as a result of only having a heavy vowel or being closed, though when we need to refer to them in particular we will describe them as *superheavy syllables*. Similarly, we leave the question of whether certain codas (e.g., sonorants) contribute weight and others do not (see Zec, 1995) to future work.

As previously noted, heavy vowels show an important alternation in French. In final open syllables, underlyingly heavy vowels are typically phonetically short (like light vowels), whereas in final closed syllables, those same vowels are phonetically long (unlike light vowels) (e.g., Côté, 2012; Goad and Prévost, 2011; Montreuil, 1995; Walker, 1984). This suggests that a different behaviour is expected for vowel weight compared to coda weight in that codas always render a syllable heavy, while underlyingly heavy vowels do not.[7] We expect that final syllables will be sensitive to the *source* of weight, heavy vowel vs. coda, whereas penults will be sensitive simply to the *presence* of a heavy syllable.

We provide example words and their expected prominence patterns in Tables 1 and 2 below. Table 1 shows that we expect prominence to shift to a heavy penult only in cases when the final syllable is closed, regardless of whether the final syllable contains an underlyingly heavy vowel. Table 2 illustrates that in penults, unlike for final syllables, we expect both codas and heavy vowels to render a syllable heavy and therefore that the syllable will attract prominence. We

[7] Nasal vowels may be an exception to this: nasal vowels can variably be diphthongised in final open syllables, suggesting that they may maintain their weight (as heavy). For the present paper, they have been grouped with the heavy oral vowels because they both show the pattern of being long and diphthongised in closed final syllables. We leave further examination of this issue for future work, but note that this predicts that, in our data, underlyingly heavy vowels in final open syllables may be slightly longer than light vowels on average due to nasal vowels not consistently showing the same shortening effect as other heavy vowels.

assume that superheavy syllables will pattern similarly to the other heavy syllable shapes in Table 1 (with the exception of the particularity that final open syllables with underlyingly heavy vowels will not protect final prominence), because the presence of a coda is sufficient for a syllable to be considered heavy and therefore to attract prominence.

Table 1: Expected prominence location based on final-syllable weight.

		Underlying vowel weight	
		Heavy	Light
Coda weight	Closed	Final prominence favoured *entente* [ɑ̃.ˈtɑ̃t] ‘agreement’	Final prominence favoured *antenne* [ɑ̃.ˈtɛn] ‘antenna’
	Open	Prominence shift favoured *hanté* [ˈɑ̃.te] ‘haunt (PTCP.M.SG)’	Prominence shift favoured *hantais* [ˈɑ̃.tɛ] ‘haunt (PST.2SG)’

Table 2: Expected prominence location based on penult weight.

		Underlying vowel weight	
		Heavy	Light
Coda weight	Closed	Prominence shift favoured *conster* [ˈkɔ̃s.te] ‘to establish’	Prominence shift favoured *copter* [ˈkɔp.te] ‘to chime’
	Open	Prominence shift favoured *conter* [ˈkɔ̃.te] ‘to recount’	Final prominence favoured *coter* [kɔ.ˈte] ‘to code’

2.2.2 Cues

The cues manipulated to signal the right edge of prosodic domains in French have been the subject of some debate, with pitch, duration and amplitude all having been discussed as possible cues. Authors typically agree upon the role played by pitch, with high and low tones being assigned to syllables to mark phrasal boundaries (Di Cristo and Hirst, 1993a, 1996; Hirst, Di Cristo, and Espesser, 1998; Jun and Fougeron, 1995, 2000, 2002; Mertens, 1987, 1993; Post, 1993). High tones are associated with prominent syllables, both when prominence shift occurs and when it does not, leading to the frequent characterisation of pitch peaks as being a common trait of prominent syllables across varieties. Previous work has additionally suggested that pitch targets are sensitive to vowel weight in French, finding that high tones are realised on the penult more often when the penult’s vowel is heavy (Thibault and Ouellet, 1996). In addition, differences in maximum pitch have been the main correlate examined in work describing

prominence shifts, so we expect that pitch contours will be affected both by prosodic domains and by weight, such that heavy syllables will be realised with higher pitch maxima.

Prominent syllables are typically longer than non-prominent ones, which has led to debate about whether it is the primary cue (Delattre, 1968; Schwab and Llisterri, 2012; Walker, 1984) or simply a cue alongside pitch (Di Cristo, 1998; Jun and Fougeron, 1995, 2000, 2002; Santiago, 2011; Vaissière, 1991). We therefore expect that duration will pattern similarly to pitch in participating in prominence shifts and, given previous work showing that Laurentian French has longer penults on average compared to Parisian French (Ouellet and Tardif, 1996), we expect that duration may be robustly used as a cue in the variety under examination.

It is worth highlighting that duration is confounded with weight: having phonologically long segments (heavy vowels) or additional segments (codas) is, of course, expected to affect rhyme durations because there is more content to be pronounced independent of weight. We expect some degree of a trade-off between syllables, however, with the prominent syllable being lengthened and the non-prominent syllable being compressed. In this scenario, a coda's effect on duration would not solely be increasing the closed syllable's duration on account of the additional segmental content. The effect of weight on duration should therefore be particularly robust, so we would not conclude that there is sufficient evidence to support our hypothesis that prominence assignment is sensitive to weight if a small change in duration is the only effect of increased weight.

We additionally consider the possibility that the cues will not pattern together, for example if the pitch peak remained on a light final syllable but the heavy penult showed a large increase in duration, in which case we would have evidence that both word-level and phrase-level prominences are signalled simultaneously, albeit on different syllables.

Finally, unlike pitch and duration, amplitude is not typically reported in acoustic studies on prominence in French. Indeed, some authors have proposed that amplitude is not a possible cue to prominence in French because it is associated with word-level – and not phrasal – prominence (Féry, 2013): languages with lexical stress use amplitude to signal word-level prominence, but not to mark phrasal prominence. However, French speakers have been shown to use increased amplitude as a cue to stress in Spanish in experimental settings, with the authors describing amplitude as a cue that is also used in native French (Féry et al., 2011). Additionally, amplitude

is manipulated for signalling prominence in Swiss French (Schwab and Llisterri, 2012), though this may be a feature unique to that variety. In view of most of the previous literature, we do not expect that amplitude will be associated with the marking of the right edge of prosodic domains. However, if prominence shift serves to highlight word-level properties, we expect that amplitude will pattern with pitch and duration and therefore be significantly affected by weight.

2.2.3 Prosodic Domains

As previously noted, the smallest domain of prominence in French is not the word, but instead, only higher phrasal domains assign pitch targets. Many terms have been used to describe the domains involved, but the smallest domain of prominence consistently groups together lexical words and their preceding syntactic dependents (e.g., Di Cristo and Hirst, 1993a, 1993b, 1996; Jun and Fougeron, 1995, 2000; see also Delais-Roussarie, 1996).[8] We follow Jun and Fougeron (1995) in calling this domain the *accentual phrase* (AP). APs are typically characterised by rising intonation at the right edge, with the high tone target normally being associated with the final syllable (Jun and Fougeron, 1995; see Kaminskaïa, 2009, 2015 for a variety of Laurentian French).

APs are combined into larger units called *intonational phrases* (IPs), which typically correspond to a sentence (Jun and Fougeron, 1995). For clarity, we will refer to any word that is at the right edge of an IP as IP-final without indicating that it is also AP-final (IP-final words are by definition AP-final), and we will therefore refer to any word that is at the end of an AP but *not* at the end of an IP as AP-final. Different kinds of IPs are associated with different pitch contours or tone targets based on the function of the sentence.[9] For example, declarative IPs are typically associated with a final low tone, while interrogative IPs are typically realised with a final high tone (Jun and Fougeron, 1995, 2000, 2002; Kaminskaïa, 2009, 2015; Martin, 2004; Post, 2000). In this paper, we focus only on declarative IPs, given the content of the text from which the data are drawn, and we will control for the prosodic domain type when examining

[8] A more detailed description of phrasing considerations will be provided in section 3.2.2 alongside examples.
[9] As we will be discussing both pitch as an acoustic cue (a phonetic measurement, here measured in semitones) and pitch targets (a phonological category), we will use *pitch* to refer to the acoustic measurements and *tone* to refer to the phonological target.

weight effects.[10] More specifically, we compare APs and IPs to each other in order to determine which cues are manipulated to mark prosodic domains in Laurentian French, and will include amplitude and duration in addition to pitch.

The pitch differences between APs and IPs in European varieties of French have been studied considerably. As alluded to above, for most varieties, the right edges of APs are marked with a final rise (LH*; following Jun and Fougeron, 1995). The H tone target is marked with an asterisk, which conventionally indicates a pitch accent. Following Gordon (2014), a pitch accent tone is crucially the only one in its phrasal domain, where it is associated to a syllable that is stressed or prominent rather than aligning with a morphosyntactic boundary. However, H* in French is often assumed to be assigned only to the phrase-final syllable, which suggests that it is a boundary tone rather than a pitch accent, consistent with its alignment with the phrase boundary.

The right edges of declarative IPs are typically marked with a low boundary tone (L%), which replaces the high (H*) of the AP's rise and makes the contour level or slightly falling (LL%). Here, the percent sign indicates that the tone is a boundary tone (and is therefore associated to a phrase edge). We expect that AP-final syllables are marked with a high tone – and therefore high pitch – that is preceded by a rise predominantly occurring on the final syllable. IP-final syllables have low pitch due to the IP-final low tone.

Though the pitch differences between APs and IPs have been studied extensively, it remains unclear whether duration and amplitude are also manipulated to distinguish between these prosodic domains. Previous work is divided about whether APs or IPs are realised with longer final durations, with APs being observed to have longer final syllables in some cases (Féry et al., 2010) and IPs being found to have longer final syllables in others (Michelas et al., 2010, though only for slow speech). This suggests that, if there is a significant difference in the degree of final lengthening based on type of domain, it is small and therefore small-scale studies are likely to

[10] The read passage contained an in-text question that was excluded from the analysis because we did not have enough data from questions to reliably examine the realisation of interrogative IPs, which are known to have a different intonational contour (e.g. Beyssade et al., 2007). The title (which was phrased as a question) was similarly not included, though it was typically pronounced with falling intonation.

find no result – or seemingly contradictory results – simply by chance.[11] The literature does not suggest an expected difference in penult durations for prosodic domains of different sizes, though it is worth noting that, in languages with lexical stress like English, the last syllable with primary stress in the phrase can be the target of final lengthening instead of the final syllable (e.g., Shattuck-Hufnagel and Turk, 1998). We expect that final syllables at the right edge of IPs are longer than those at the right edge of APs, but for this difference to be small based on mixed results in previous work.

Finally, the use of amplitude to mark different types of prosodic domains in Laurentian French, if it is manipulated at all, is not yet known. Based on work on other dialects and on typology, we expect that amplitude will not be manipulated to distinguish between APs and IPs for any syllable.

2.2.4 Summary

Based on the background literature, we expect that French – and Laurentian French in particular – predominantly uses pitch and duration to mark prosodic domains. The classification of French within the typology of prominence systems is at the centre of some debate (e.g., Vaissière and Michaud, 2006, who describe French as a "non-tone, non-stress language"): authors differ in whether they describe final prominence in French as a stress (e.g., Cutler, 2005; Schwab and Llisterri, 2012), as a pitch accent (e.g., Jun and Fougeron, 1995; cf. also Rossi 1980), or whether they do not formally categorise it, and there has additionally been some suggestion that the prosodic system is – or recently was – in flux (e.g., Fónagy, 1980).

We additionally expect French to exhibit weight sensitivity, with both long vowels and closed syllables being heavy and therefore attracting prominence. These heavy syllables are expected to be marked with increased rhyme durations and higher amplitudes. Furthermore, we expect the tone target (the H* in the AP's LH*) to shift inwards towards heavy penults and away from open final syllables, leading to higher pitch maxima for penults compared to final syllables. Determining the patterns found in these contexts would allow us to better situate French within

[11] It could additionally be that other factors confound the results of previous studies for duration; for example, the presence of a following pause could be associated with reducing an IP's final lengthening, since the cue is no longer needed. We will not address the availability or use of other cues in this paper.

the typology of prominence systems. This is important in part because it could allow us to better interpret the results provided by French speakers in perceptual studies, where participants often pattern differently than expected based on the description of French as a language with prominence strictly on the final syllable (e.g., Frost, 2011; Li et al., 2017).

Figure 1 illustrates the pitch profiles we expect by showing the idealised pitch contours at the right edge of the accentual phrase for unshifted (top row) and shifted (bottom row) cases of APs (left) and IPs (right). The values in the figure were chosen to produce curves with shapes consistent with our hypothesis in (1) that prominence assignment is sensitive to weight. In APs that are not IP-final, we predict a rising contour with the H* aligned with the final syllable if shift does not occur, but with the penult if shift does occur. In IP-final contexts, however, we expect to get no high tone if shift does not occur, since the H* in the final syllable gets replaced by the IP's L% boundary tone, assuming a declarative sentence. When shift does occur, the AP's H* tone would be expected to move inwards to the penult and therefore the phrase would end in a rise-fall because the IP's L% tone would be aligned with the boundary and would not replace the AP's H* that is assigned to the previous syllable.

Figure 1: Predicted pitch contours at the end of APs depending on whether the AP-final word is also IP-final (right panels) or not (left panels), and whether the word undergoes prominence shift (bottom panels) or not (top panels).

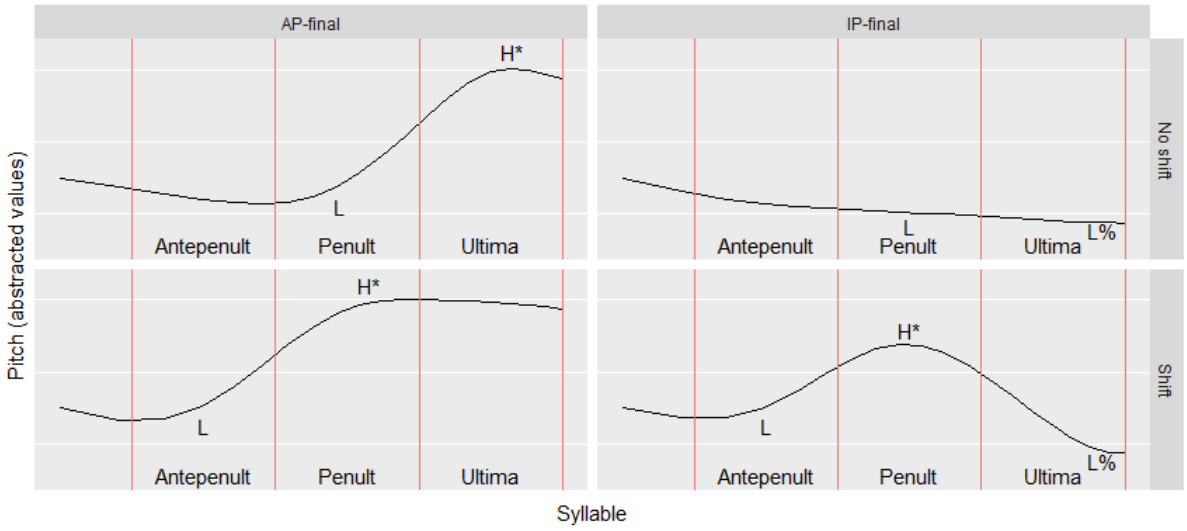

The next section will address the data preparation, coding and statistical analysis that were used in order to test the patterns we predict for the realisation of prominence in the Laurentian French data under examination.

2.3 Methods

The goal of this study is to quantitatively test whether phonologically heavy syllables attract prominence in Laurentian French. Section 2.3.1 discusses the choices of corpus and of region under examination. In section 2.3.2, we turn to the procedure for data extraction and processing, in particular noting decisions about syllabification, the coding of prosodic boundaries, and the extraction of acoustic cues. Section 2.3.3 then discusses the mixed-effects logistic regression models, and how factors were coded for those models. Finally, section 2.3.4 presents the predictions for those factors.

2.3.1 Corpus and Speakers

As mentioned earlier, this study draws its data from the Laurentian sub-corpus (Côté, 2014, 2015) of the *Phonologie du français contemporain* corpus (PFC; Durand et al., 2002, 2009; http://www.projet-pfc.net/), which is a large-scale project seeking to provide corpus data from varieties of French around the world. During corpus collection, speakers from each survey location take part in four tasks: (1) reading a list of words designed to elicit regionally variable word-internal phonological phenomena in more attentive speech, (2) reading a short passage intended to examine sandhi processes and cross-word variability alongside word-internal phenomena in less attentive read speech, (3) having a semi-directed conversation lasting 20 to 30 minutes, and (4) having an informal conversation lasting 20 to 30 minutes.

In this study, we focus on the read passage (provided along with its translation into English in Appendix A). This was to ensure that the tokens across speakers are more comparable and that the phrasing is relatively fixed, since the speakers generally formed prosodic domains based on the punctuation provided in the text. Additionally, speech rate and register tend to be more consistent throughout a read passage than in spontaneous speech, which further increases comparability within and across speakers. Finally, work comparing dialect differences in French prosody has shown that read speech prosody is more similar across dialects than spontaneous speech prosody is (Simon, 2003), which means that the results obtained in this study are more likely comparable to those that would be found for other dialects.

As previously noted, we chose to examine Laurentian French because it has conserved a large number of vowel length contrasts in comparison to other dialects, providing more opportunities for weight effects to be observed. For example, we find that the mid-vowel pairs /e-ɛ/, /ø-œ/ and /o-ɔ/, all listed with the heavy vowel first, are still robustly distinguished (e.g., Côté, 2012; Lamontagne, 2014). Examples of this phonemic contrast can be seen in pairs like *fée* /fe/ 'fairy' and *fait* /fɛ/ 'fact', and like *côte* /kot/ 'hill' and *cote* /kɔt/ 'code'. There is additionally a long counterpart to /ɛ/, transcribed as /ɛː/ and frequently realised as diphthongised (Côté, 2012), which we see contrasted in pairs like *fête* /fɛːt/ 'party' and *faite* /fɛt/ 'done (FEM)'. Furthermore, the variety has conserved a contrast in the low vowels, as seen in *pâte* /pɑt/ 'dough' and *patte* /pat/ 'paw', the former being underlyingly heavy. The presence of these additional vowel contrasts make it so that variation in vowel weight is more common in Laurentian French, and more pronounced as a result of diphthongisation, making this variety an optimal starting point for testing the prosodic effects of syllable weight in French.

More specifically, we examine the Laurentian French variety spoken in Chicoutimi, Quebec, which is located 200 km north of Quebec City. This area was selected for demographic reasons: we wanted to ensure that there would be limited contact with other languages and dialects in case this could influence the results within or across individuals, thereby minimising the risk of substrate influences or other contact effects on prominence. The Saguenay area, which includes Chicoutimi, is optimal for this, as census data (Statistics Canada, 2012) show: 98.3% of inhabitants report speaking French as a native language – higher than the Quebec-wide average of 78.1% –, and 98.9% speak only French at home. The frequency of French usage at home increases to 99.9% if all inhabitants who speak French at home on a regular basis are included.

The rate of bilingualism, including French-English bilingualism, is also relatively low. Inhabitants aged 20 to 44 are the most likely to be bilingual (31.8%), with those aged between 45 and 64 and those over 65 having bilingualism rates of about 17% and below 13%, respectively (Statistics Canada, 2012). These figures reflect lower bilingualism than for the province as a whole by about 20 percentage points for each age group. The Saguenay area additionally sees relatively little immigration, both from outside of the country and from elsewhere in the province or country, meaning that contact with other varieties of French and contact with other languages, including English, is unlikely to affect the results.

The data in this study come from 11 native French speakers who were born in and who grew up in Chicoutimi, with speakers spanning three generations and being relatively well-balanced for sex. The speakers' demographic factors – sex, birth year and age at time of testing – are presented in Table 3. We can see that the speakers are not perfectly balanced for these demographic factors, but previous research on French suggests that there are no gender effects with regards to the placement of prosodic tone targets (Avanzi et al., 2011). In all cases, the surveyed speakers were fluent readers.

Table 3: Speakers in the Chicoutimi survey and their demographic factors (year of birth, age)

	Sex	
Age Group	Male	Female
Young adult	fv1 (1984, 22) pl1 (1983, 23)	cl1 (1982, 24) jv1 (1979, 27) mb1 (1985, 21)
Middle-aged	db1 (1954, 52) pt1 (1965, 41)	gm1 (1958, 48) ma1 (1953, 53)
Older	rt1 (1934, 72)	gt1 (1932, 74)

2.3.2 Procedure

2.3.2.1 Alignment and Syllabification

The data were automatically processed to obtain acoustic measurements and lexical information. The PFC recordings were first forced-aligned using an aligner that was trained on Laurentian French and that uses speaker-adapted word-internal triphone models to maximise accuracy (Milne, 2014). From there, a Praat (Boersma and Weenink, 2016) script created rhyme and syllable tiers using an onset-maximisation algorithm that checked whether a consonant cluster formed an allowable onset. Word-internal /VsCV/ sequences were syllabified such that the /s/ was in coda, which reflects distributional tendencies in French: for example, /s/ does not occur in word-medial codas after a nasal vowel (historically vowel-nasal sequences) morpheme-internally, suggesting that /s/ is sensitive to constraints on coda, unlike the first consonant of a medial cluster with rising sonority. In all cases, the syllabifications used in this analysis are

based on realised forms and not possible underlying ones, as coded by the forced aligner; reduced clusters and omitted schwas are not reconstructed, so that a word like *mettre* /mɛtʁ/ 'to put' was considered as having one syllable if the possible final schwa was not realised (e.g., [mɛt], [mɛtʁ]), but two if it was (e.g., [mɛtʁə]).[12]

2.3.2.2 Identifying Prosodic Domains

As we specifically examine prominence shift to the penult, only words with at least two realised syllables were extracted.[13] Accordingly, we also only examine the last two syllables of each target word, regardless of how many syllables are in the word. We additionally restricted the words analysed to those that are at the end of an accentual phrase given that this is the smallest prosodic domain described as assigning pitch targets (e.g., Jun and Fougeron, 1995, 2000, 2002). Since accentual phrases are always contained within intonational phrases and intonational phrases always end in an accentual phrase, this also means that all IP-final words of at least two syllables have been extracted. This was done manually based on the text – identical for all speakers – to ensure that the coding was not biased by the presence or absence of cues that we would consciously associate with the right edge of a domain, which could affect the results.

The contexts included are shown in (2)-(4),[14] where we see examples of AP-final tokens that are taken directly from the text analysed (and are deemed to be AP-final, following the criteria from Jun and Fougeron, 1995). These right edge boundaries, indicated using pipes ("|") in the examples, correspond to syntactic junctures and are further supported by work that examines the prosodic domains relevant to phonological processes in French, like liaison (Hannahs, 1995).

[12] The words were also manually coded for whether the penult was the last syllable of the base, so that words like *gouvernement* 'government' and *protéger* 'to protect' were marked as having a base-final penult (cf. *gouverne* 'governs' and *protège* 'protects'). Given the limited data, we cannot speak to morphological effects directly with confidence and leave this question to future work, but we control for them through the factor and by-word random effects in the statistical model used in the current study (see further section 3.3).

[13] Jun and Fougeron (1995) found that prominence can shift to the antepenult as well. However, since it represents a small proportion of their realisations and since our data primarily consist of two-syllable words, we focus only on shifting prominence to the penult. We also leave for future work prominence shift to non-final words in an AP that ends with a monosyllabic word; under the conventional description of prominence in French and in our final analysis of phrasal prominence, it is expected that shifting prominence to a non-final word would be possible.

[14] The translations correspond to the senses of the words as they appear in the text, taking into account portions of the text not included in the abridged examples. Where example passages also include prosodic boundaries of types other than the one being illustrated, only the boundaries targeted for the example have been marked. Nouns followed by an adjective have been excluded because they can be realised in the same or separate APs (Post, 2003), as shown in section 1.

When punctuation was present, as in (3) and (4), we treated the prosodic context as being distinct from when there was no punctuation present, leaving us with three groups that were included in the statistical models: AP-final tokens not followed by punctuation, AP-final tokens followed by a comma, and IP-final tokens (AP-final tokens followed by a period). Only two of these groups are under present focus: AP-final tokens without punctuation, which, as mentioned earlier, we refer to as AP-final, and AP-final tokens followed by a period, which we refer to as IP-final.[15] In total, 1368 tokens were included, meaning that there were 2736 syllables measured.

(2) Contexts coded as AP-final without punctuation:
   a. After a noun
      Ex: ... *le gouvernement* | *prend contact* | *avec la préfecture* | ...
          ... the government | makes contact | with the prefecture | ...
   b. After a verb that is not an auxiliary
      Ex: ... *le gouvernement ... et s'assure* | *que* ...
          ... the government ... and ensures | that ...
   c. After a post-nominal adjective
      Ex: *La côte escarpée* | ...
          The steep hill | ...
   d. After a post-verbal adverb
      Ex: ... *qui tournaient toujours* | ...
          ... that always pivoted | ...

(3) Contexts coded as AP-final with punctuation:
   a. After any word followed by a comma
      Ex: *Le hasard* | *tout bêtement,* | *car le Premier Ministre,* | ...
          Happenstance, | quite frankly, | since the Prime Minister, | ...

(4) Contexts coded as IP-final:
   a. After any word followed by a period
      Ex: ... *depuis les élections.* | ...
          ... since the elections. | ...

[15] AP-final tokens followed by a comma were excluded from the analysis for reasons of space. In previous work (Lamontagne et al., 2017), we discuss the results for this domain and their implications for the assignment of domain-marking tones.

2.3.2.3 Acoustic Measurements

For each syllable that was included in the analysis, a Praat script extracted the rhyme durations,[16] as well as the syllable's maximum pitch and its maximum amplitude, based on the cues discussed by Gordon (2014) in his typological work and those examined for French by Jun and Fougeron (1995, 2000, 2002).[17] We focus our analysis on maximum pitch (discussed further in section 2.3.4.1) because it corresponds to the phrase-final high tone that is interpreted to be the acoustic target for AP prominence (e.g., Jun and Fougeron, 1995, 2000, 2002). Figure 2 illustrates an example of a token without shifted prominence (speaker cqbcl1), while Figure 3 shows a token with prominence shift (speaker cqbfv1). In these tokens we can see the pitch maximum is highest in the final syllable and in the penult, respectively. In both cases, the token was AP-final and was followed by segments with lower pitch at the beginning of the next AP within the same IP.

It is worth noting that with this method we do not directly know whether a contour is rising or falling, but the height of the maximum pitches for the two syllables is suggestive of the contours present and directly reflects the component of prominence that interests us with respect to pitch (the tone target). In Figures 2 and 3, for instance, the syllable with the highest maximum pitch shows a pitch rise leading up to that maximum pitch. In a case where the token is also IP-final, a fall or low plateau is present, and therefore the contour can be inferred effectively from the values.

[16] We additionally extracted vowel durations to examine which part of the rhyme is particularly affected by lengthening because in discussion of length alternations in French, it is specifically the vowel that is said to lengthen (Walker, 1984). However, we will not report on the vowel duration data as the patterns under present focus for the rhyme and vowel were the same.

[17] Earlier analyses, such as those presented in Lamontagne et al. (2018), included results for pitch range. However, as those results closely mirror the ones found for maximum pitch and less directly reflect the high-tone target of the AP, we have not included them in the current paper for conciseness. It is worth noting that the maximum pitch measure shows compressed results because the value is always obtained from the highest point: a rise beginning in an earlier syllable will show a smaller difference between syllables than a rise that entirely takes place in a single syllable, and the lowest point of a fall will not be captured by the measure.

Figure 2: Le premier ministre ‘the prime minister’ produced by cqbcl1, in which prominence shift did not occur. The final cluster was reduced such that [s] was the end of ministre, with the alignment otherwise showing the phonemic transcription for each phoneme.

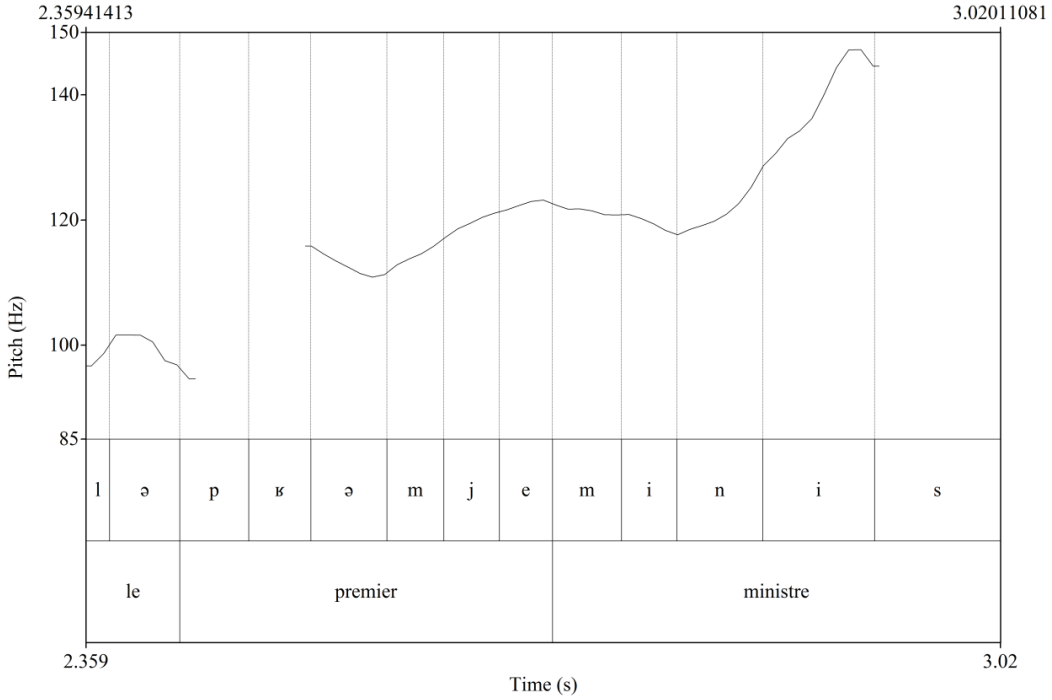

Figure 3: (Ont) eu tendance ‘(had) a tendency’ produced by cqbfv1, in which the pitch maximum was realised on the penult.

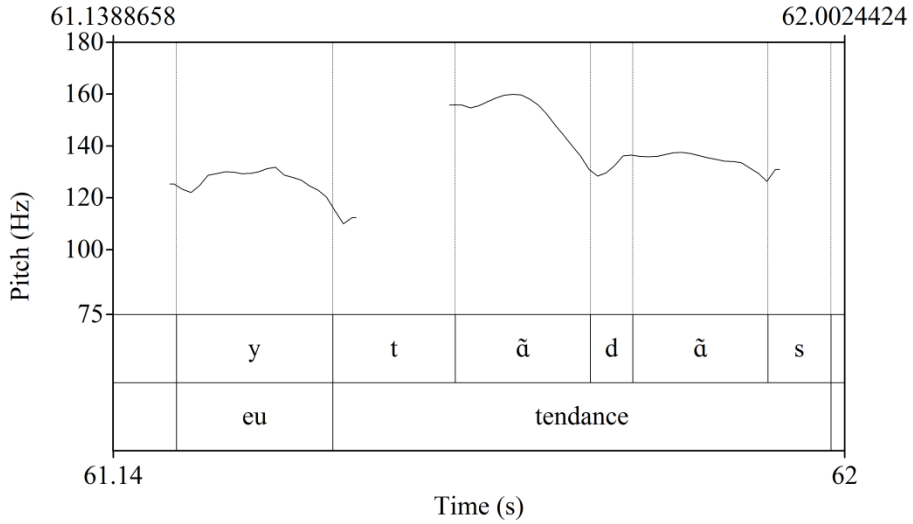

Maximum amplitude was preferred over mean amplitude because the mean is more likely to be affected by the segments present in a given syllable and because a shorter vowel would be expected to be at its maximum amplitude for a shorter period of time, thereby reducing the mean value without necessarily reflecting a lower amplitude target. Using maximum amplitude also

meant that we could reliably measure through the rhyme instead of limiting ourselves to measuring the vowel, which could have resulted in not including the point with the greatest amplitude if, for example, there was a sonorant consonant in the coda that was higher in amplitude than the immediately preceding vowel.

2.3.3 Models

The acoustic cue realisations of the 2736 targeted syllables were analysed using the lme4 package (Bates et al., 2015) in R (R Development Core Team, 2015) to compute mixed-effects linear regressions with speaker and word as random intercepts and with by-speaker random slopes for all fixed-effect factors and interactions. No by-word random slopes were included because the predictors only varied by speaker. We discuss the results of three models, one for each acoustic cue (maximum pitch, rhyme duration, and maximum rhyme amplitude). We plotted the residuals and found them to be approximately normally distributed, and we verified the correlation matrices to ensure the assumption of multicollinearity was not violated.[18]

The models take as their dependent variables not the raw acoustic measurements, but instead, the difference between the last two syllables' values, which yields a relative value to provide some normalisation for the context, with the use of random slopes and random intercepts simulating normalisation procedures to remove inter-speaker differences (Drager and Hay, 2012). For clarity, we will refer to the relative value as *RV* throughout the rest of this analysis. The formula for the RV, presented for each cue in (3), involves subtracting the final syllable's value from the penultimate syllable's value. This transformation was chosen because it provides an interpretable value: an RV greater than 0 indicates that the penult has a higher cue measurement, while an RV below 0 indicates that the final syllable has a higher value. The further from 0 that the RV is, the larger the difference between the two syllables is.

[18] While the factors are not too confounded for testing, there are distributional asymmetries that mean the data are skewed towards having certain phonemic content in some contexts. For instance, while /e/ is common word-finally, it does not occur in closed final syllables. As a result, while we have numerous tokens of heavy vowels in closed final syllables, the heavy vowels are only a subset of those found elsewhere (e.g. /o/ and nasal vowels still occur in this context). Similarly, /ɔ/ is absent word-finally, meaning that it is amongst the light vowels in closed final syllables, but it is not included in the light vowels found in final open syllables (which include /ɛ/, for example).

We use subtraction instead of division when calculating the RV for two reasons. First, the cues were already log-transformed either as a preliminary step (for duration) or as a result of their units (decibels for amplitude and semitones for pitch), meaning that subtraction is equivalent to dividing the cue values, following the laws of logarithms. Second, this allows for a more intuitive interpretation of the RVs: not only does the RV's sign indicate which syllable has a greater cue measurement, but the RV can be directly interpreted as being the size of the effect – an RV of 1 for the pitch maximum means that the penult's pitch maximum is one semitone higher than the one for the final syllable, for example. Because our models look at RVs, the model considered 1368 data points (one per word rather than one per syllable). Four tokens were excluded from the maximum pitch and maximum amplitude models due to excessive devoicing of a high vowel making those measurements unreliable.

(3) Formula for calculating RVs:

$RV_{cue} = measurement_{penult,cue} - measurement_{final,cue}$

Figure 4 illustrates how RVs relate to the cue measurements using hypothetical values generated through the *rnorm* function in R.[19] In the panel on the left, we provide an example of what the hypothetical results for penult weight could be – penults having a higher value when heavy compared to when light. In the panel on the right, we see the RVs based on the hypothetical measurements: the RV is positive when the penult vowel is heavy because the penult's value is larger than the final syllable's value; the RV is negative when the penult vowel

[19] 500 tokens were generated for each combination of the penult and final syllables being light or heavy. Prominent syllables were given a mean of 25, while non-prominent syllables were given a mean of 15. If a final syllable was heavy (ie. the final syllable is expected to preserve prominence) or the penult was light (ie. the penult is expected to attract prominence), the final syllable was treated as prominent and the penult was treated as non-prominent. To reflect the hypothesis that heavy penults will optionally attract prominence from light final vowels, the penult was treated as prominent (mean value of 25) in half of the cases where the penult syllable was heavy, with the final syllable therefore not being treated as prominent (mean value of 15). In all other cases, the penults were treated as non-prominent (mean value of 15), with the final syllable instead being prominent (mean value of 25). If the two syllables had equal weight and therefore more variation in prominence may occur, the standard deviation for both syllables was set to 7, while it was set to 3 if the syllables differed in weight. While we do not believe the differences in means and variation are this straightforward in reality, it provides a straightforward set of values for visual inspection of potential patterns and illustrates that even a relatively clear effect is not as evident in plots. Figure 4 and the associated discussion only include the tokens that were generated to include prominence shift (ie. penults always attract prominence when the penult is heavy and the final syllable is light), but Figure 5 and the associated discussion in section 3.4.5 introduce the additional tokens.

is light because the final syllable's value is larger than the penult's value.[20] In the statistical analysis, the model coefficients would reflect the penult attracting prominence by being a positive value proportional to the difference in RVs between the light-vowel context and the heavy-vowel context.

Figure 4: Syllable weight values and their associated RVs using hypothetical data.

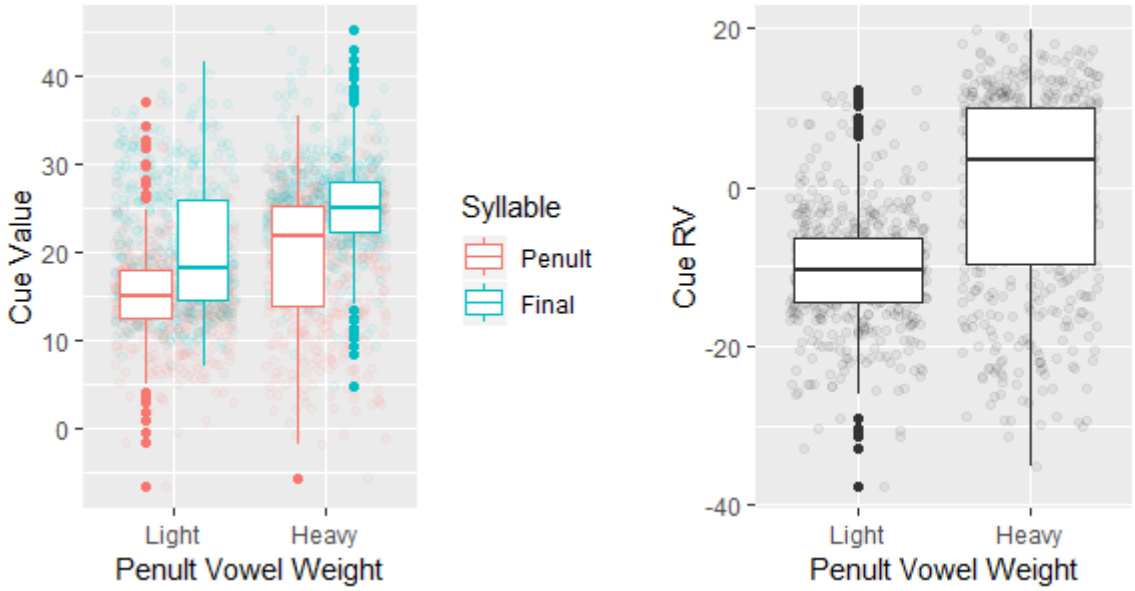

Returning to the statistical analysis, the models all include the same fixed and random effects to ensure that they are maximally comparable. With the exception of the prosodic domain, all factors are binary and were therefore rescaled by two standard deviations and centred for better comparability with other analyses. The prosodic domain is a ternary factor (AP with no punctuation vs. AP with comma vs. IP) and was Helmert-coded so that the first prosodic domain factor in the model provides the difference between APs with no punctuation and IPs, as this is the domain comparison we focus on in the analysis. In all cases, except when specifically discussing intonational pitch targets, the following directions of effect are interpreted as an increase in prominence: higher pitch and amplitude maxima and longer durations. The

[20] For maximum pitch, the RV will often be near 0 in the case of tone targets retracted to the penult because the final syllable will not have a separate tone target, as shown in Figure 1, and therefore the final syllable's pitch will be similar to the penult's pitch.

coefficients in the model can be interpreted as the size of the change in acoustic cue measurements.

We included one additional factor in the model that does not directly relate to the predictions within the scope of this paper: the morphological structure. A larger – and more morphologically diverse – dataset would be required to test the effects of morphological structure in detail since many properties of individual morphemes could play a role (e.g., syllable shape, phonological size, status as derivational or inflectional). However, preliminary data exploration suggested that we would need to control for morphological structure in our models, so a manually coded factor that identified whether the penult was base-final or not was included. We will point out where this factor was crucial in our description of the results, but not treat it as a focus given the limitations of the dataset.

2.3.4 Predictions

In this section, we discuss predictions that stem from our hypothesis that French prominence is sensitive to weight. Anticipating the order of presentation in the results section, we begin by providing our predictions relating to prosodic domains in section 2.3.4.1. We then turn to coda weight and vowel weight, in sections 2.3.4.2 and 2.3.4.3 respectively. In section 2.3.4.3, we also provide our predictions for how vowel weight and coda weight interact in the final syllable. Finally, in section 2.3.4.4, we discuss how our predictions differ slightly between penultimate and final syllables.

2.3.4.1 Prosodic Domains

AP-final syllables are generally expected to be marked with an LH* sequence, which means that the penult will have a lower maximum pitch (marked with a low tone) than the final syllable (marked with a high tone), yielding a negative RV for maximum pitch. We also expect IP-final syllables to be marked with an LL% or H*L% contour, which means that the RV will be near or above 0 since the final syllable will be marked with a low tone. These expectations lead to the prediction that the maximum pitch RVs will be significantly higher in IPs (where the final syllable has a low tone) compared to in APs (where the final syllable has a high tone).

Given the results of previous studies for final lengthening in different domains, we expect that the duration RV will be higher in APs than in IPs because the final syllable of IPs will be subject to greater phrase-final lengthening. However, this difference is likely to be small based on the mixed results in the literature, and this study will not have the statistical power required to confidently refute that a small effect exists and therefore will need to be compared alongside other work in meta-analyses. These expectations lead to the prediction that we will find any difference in duration resulting from the type of prosodic domain.

Finally, we do not expect that there will be a significant difference between APs and IPs with respect to amplitude RV because amplitude is not expected to be used as a cue to phrasal prominence, though again we will not be able to refute the presence of a small effect if no statistically significant effect is found. We therefore do not predict any difference in amplitude resulting from the type of prosodic domain present.

We summarise our predictions for the marking of prosodic domains as Prediction 1 in (5):

(5) **Prediction 1**: We predict that IPs will have lower-pitched final-syllable rhymes than APs will, but we do not predict that there will be a significant difference in duration or amplitude between APs and IPs.

2.3.4.2 Coda Weight

The first possible source of weight that we discuss is codas, where closed syllables are heavy and open syllables (with short vowels) are light. We expect that syllables with codas will have greater prominence compared to those without.[21] As a result, our expectations for coda weight are straightforward: closed penults will attract prominence, while final closed syllables will preserve prominence. These expectations lead to the prediction that closed syllables will have higher maximum pitch RVs, longer rhyme duration RVs, and higher maximum amplitude RVs.

[21] French has a process whereby the vowel in a final syllable is longer when the syllable is closed by /v z ʒ ʁ vʁ/ than when the syllable is either closed by another consonant or open (e.g., Côté, 2012; Walker, 1984). This length difference is particularly evident in high vowels in Laurentian French since laxing and lengthening, both of which occur in closed syllables, are mutually exclusive: compare *vite* [vɪt] ‘quick’ and *vivent* [vi:v] ‘live (3PL)’ (also *vie* [vi] ‘life’). Because this lengthening process was shown not to have an effect on the realisation of prominence in earlier statistical models (see Lamontagne et al., 2018), this factor was not included in the models discussed here.

The RVs will therefore be higher if the penult is closed rather than open, and lower if the final syllable is closed rather than open.

We summarise this second prediction in (6):

(6) **Prediction 2**: We predict that closed syllables will have significantly higher values compared to open ones for all acoustic cues signalling prominence (maximum pitch, duration, and amplitude).

2.3.4.3 Vowel Weight

Before turning to our predictions for vowel weight, we discuss how we partitioned the French vowel space into length classes. As previously noted, we modelled vowel weight as binary. Following Walker (1984) and Côté (2012), we classified vowels into the two categories in Table 4, which we label as light and heavy. Vowels that are classified as heavy are those that surface as long in closed final syllables, thereby revealing their underlying weight. Although one vowel in this category, /e/, does not occur in this context in native French words, it was classed as heavy based on its phonological similarity to /ø o/, on its phonological opposition to the unambiguously short /ɛ/, and on how it surfaces as long – and often diphthongised – in integrated borrowings from English like *break* or *tape* (Côté, 2012; Lamontagne, 2014). For both syllables under analysis, the binary factor for vowel weight (light, heavy) was centred and rescaled by two standard deviations.

Table 4: Vowels in Laurentian French and their inherent weights.

Light vowels	Heavy oral vowels	Heavy nasal vowels
i y u		
	e ø o	
ɛ œ ə ɔ	ɛː	ɛ̃ œ̃ ɔ̃
a	ɑ	ɑ̃

We expect that heavy vowels will attract prominence and therefore that (a) the final high tone will be attracted to heavy vowels, (b) that heavy vowels will have significantly longer duration, and (c) that heavy vowels will be marked with increased amplitude. Combining our expectations

that final syllables must be closed to count as heavy and that vowel weight contributes to weight in final closed syllables, we expect that final syllables that both are closed and have a heavy vowel will preserve prominence, but that final open syllables will pattern as light even if they contain an underlyingly heavy vowel. These expectations lead to the prediction that heavy penult vowels will be associated with higher maximum pitch RVs, higher duration RVs, and higher maximum amplitude RVs, but we do not predict that an effect will be found in final open syllables.

This prediction is summarised in (7):

(7) **Prediction 3**: We predict that syllables containing a heavy vowel will have significantly higher values for all acoustic cues (maximum pitch, duration, and maximum amplitude) compared to syllables containing a light vowel, except in the case of final open syllables where we do not predict that a significant effect will be found.

2.3.4.5 Differences between Syllables

In the discussion thus far, we have treated the two syllables under focus as equally capable of hosting prominence. This, though, is not consistent with the literature where final syllables are standardly considered to be the default position for prominence assignment in French. In view of this, we must modulate our predictions to ensure that the prominence-retaining properties of final syllables will have a greater effect than the prominence-attracting properties of penult syllables. This may manifest itself not only in the relative sizes of the predicted coefficients (larger for factors relating to final syllables than for those relating to penults), but also in the distributions themselves.

Consider again Figure 4, where we presented hypothetical results for weight effects to illustrate how the RVs relate to the acoustic values. Examining the predicted role of penult vowel weight, we see that the hypothetical distributions are entirely separate – as though the profiles of the final syllables do not have any effect. If we revise the simulated data to better match the hypothetical patterns in Table 2, this would more closely match the distributions in Figure 4. For this figure, we treated the penult as prominent when the penult was heavy and the final syllable

was light, but the final syllable as prominent when this was not the case (i.e., heavy final syllables retain prominence, and light penults do not attract it).

We can see from a comparison of Figures 4 and 5 that this has a fairly marked effect on the appearance of the distributions in plots: now even when the penult is heavy, the distribution suggests frequent final-syllable prominence because it is centred closer to zero, and often has a mean below zero. The plots in the results section present the acoustic measurements in the same format as in the left panels of Figures 4 and 5 (with the syllable values rather than RVs) to make the patterns easier to interpret visually, but the models take as input the RVs (like what is shown in the right panels of Figures 4 and 5). With these hypothetical data, a fixed-effects linear regression model that includes the weight of both syllables and the interaction between those weights gives a coefficient for weight of 2.425 (positive), after centring and rescaling the vowel weight factor. This means that heavy penults increase the mean difference in cue values by about that number relative to the estimated mean (and therefore by twice that number relative to the penult being light, since the mean is centered between the two values).

Figure 5: Hypothetical results and their associated RVs.

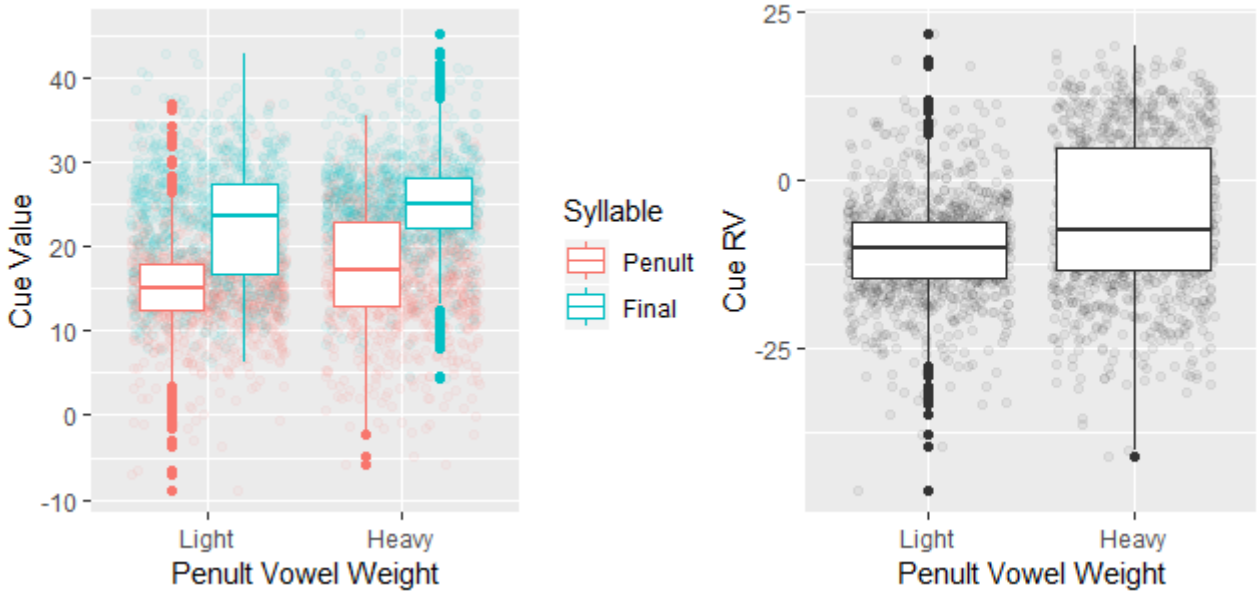

2.4 Results

In this section, we discuss the results of the statistical models. We present the findings thematically – based on the main factors of interest – to directly compare the effect of each factor on the acoustic cues. We begin with the results for prosodic domains in section 2.4.1, discussing how the marking of the right edge of IPs compares to the marking of the right edge of other APs. We continue in sections 2.4.2 and 2.4.3 with the results for coda weight and vowel weight, respectively. Section 2.4.4 then examines the interaction between coda and vowel weight in the final syllable – in other words, it allows for a comparison of final syllables that are underlyingly light (open with light vowel), heavy (open with heavy vowel or closed with light vowel) and superheavy (closed with heavy vowel). The model results themselves are provided in Appendices B through E, and for ease of reference all figures in this section will follow the same layout, where panel A shows maximum pitch, panel B shows rhyme duration, and panel C shows maximum amplitude.

2.4.1 Prosodic Domains

Prediction 1 stated that IPs would have higher pitch maximum RVs than APs because the IP-final syllable receives a low tone. We find that IPs have considerably higher values (β=1.1817, $p<0.0001$), which is consistent with IPs having a low tone rather than a high tone in the final syllable. We see this in panel A of Figure 6, where we additionally note that this difference mainly appears to be a result of the final syllable's pitch changing, consistent with these phrase types having different final-syllable targets, but not different tone targets for the penult.

Regarding duration RVs, we predicted lower values in IPs than in APs because we expected to find more final lengthening in IPs than in APs. Consistent with this prediction and as shown in panel B of Figure 6, we find a small but significant difference whereby IP-final syllables are proportionally longer than AP-final ones (β=-0.1068, $p=0.0480$). However, this result should be tested in future work, as the effect is small and only barely reaches the threshold for significance. Additionally, we note that the effect is too small to counteract the intercept (β=-0.6097, $p<0.0001$), meaning that overall final syllables are longer than penults unless other factors (such as weight, to be discussed below) lengthen the penult or shorten the final syllable.

As for amplitude, we predicted that no significant difference between APs and IPs would be found because amplitude is not expected to be manipulated to signal boundaries. Contrary to this prediction, however, we find that IPs have significantly higher RVs for maximum amplitude (β=1.9705, p=0.0011), which indicates that IP-final syllables have much lower amplitude than AP-final ones. Panel C of Figure 6 illustrates this interpretation, where we see that both the penult and final syllables have lower amplitude in IPs than in APs, with final syllables showing the largest decrease.

Overall, these results are consistent with (a) APs being marked with a rising pitch contour and (b) IPs being marked with low final pitch. We additionally find evidence that IP-final syllables are longer than AP-final ones and that IPs have lower final amplitude. These results suggest that Laurentian French patterns like other French varieties in terms of the tone targets used to mark phrasal domains. With the tone for APs established, we turn to the results for weight factors.

Figure 6: Results for domain type.

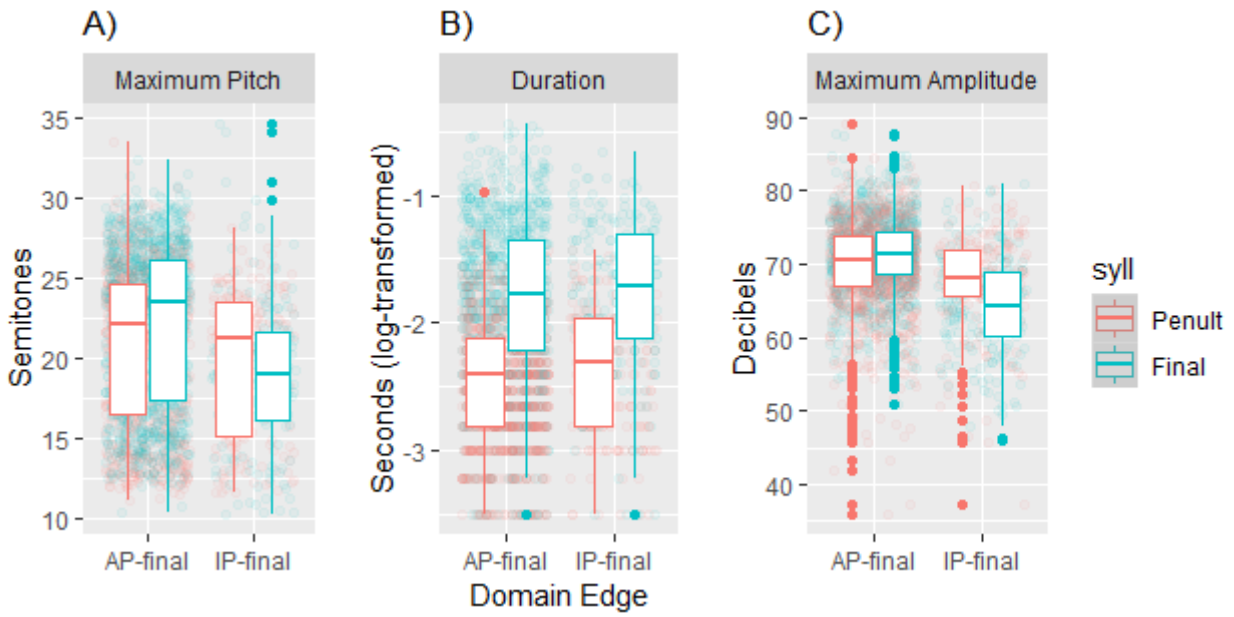

2.4.2 Coda Weight

Prediction 2 stated that the RV for all acoustic cues would be higher in closed syllables, following from our hypothesis that closed syllables will attract prominence. We begin by

presenting our results for penult coda weight in section 2.4.2.1, and then continue to final-syllable coda weight in section 2.4.2.2.

2.4.2.1 Penult Coda Weight

We predicted that closed penults would have higher pitch maxima because heavy penults attract the H* tone from the final syllable. Our models support this prediction (β=1.1533, p<0.0001), but an examination of panel A in Figure 7 suggests otherwise. Based on model comparisons, the prediction of greater penult pitch maxima is borne out so long as we control for morphological structure and for whether the final syllable is closed (which our statistical model does). However, based on additional data inspection and model comparisons, this result appears to be consistent whether the penult is base-final or not, but skewed data proportions mask this result in the figure, particularly because morphological structure has additional effects on the realised pitch contours.

Figure 7: Results for penult coda weight.

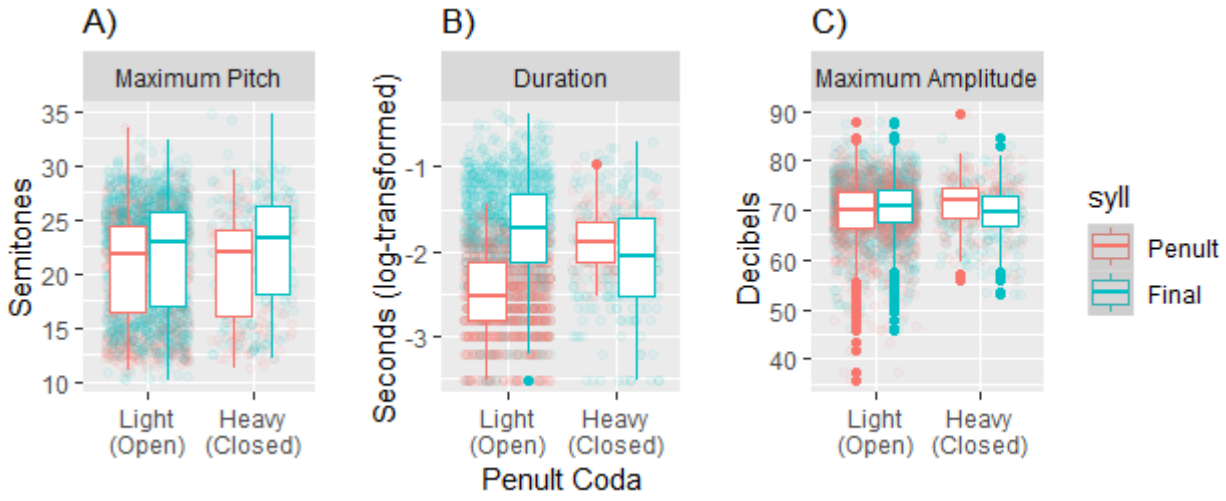

We predicted that closed penults would have significantly longer rhymes not only because of the additional segment, but because the penult optimally attracts prominence. Closed penults have significantly longer rhymes (β=0.6992, p<0.0001), and panel B of Figure 7 shows the expected trade-off between syllables, such that the penult being heavy not only increases the penult's rhyme duration, but also appears to decrease the final syllable's rhyme duration. This

relationship between the durations of the two final syllables allows us to infer that the increase in relative duration is not simply the result of there being an additional segment in the penult rhyme with no influence on prominence assignment. Finally, as shown in panel C of Figure 7, we also find that closed penults have higher amplitude RVs (β=2.0582, p<0.0084).

In short, these results show that words with closed penults, like /mɛrsi/ *merci* 'mercy', more often have higher pitch, higher amplitude and longer duration in the penult than words with open penults, like /mɛsi/ *messie* 'saviour', consistent with Prediction 2. Based on our hypothesis that French prominence exhibits weight sensitivity, we expect that closed final syllables will similarly have higher values for these acoustic cues. The next section provides the results for this prediction.

2.4.2.2 Final Syllable Coda Weight

Turning to coda weight in final syllables, we predicted that the final syllable being closed would be associated with that syllable being highly likely to preserve prominence and therefore to be realised with a high tone. We see from panel A of Figure 8 that this does have an effect (β=-0.1440, p=0.0139), though the effect is larger when controlling for other factors (like in our statistical models) than when not doing so (as in Figure 8).

Also in line with our prediction that closed syllables are heavy and therefore should be longer, closed syllables have significantly longer relative rhyme durations. While the increase in relative duration for the penult is quite large, we find an even larger effect for final syllables (β=-0.9202, p<0.0001), as shown in panel B of Figure 8, suggesting that this effect is not solely the result of adding segments.

Finally, we find much lower RVs for amplitude when the final syllable is heavy (β=-3.1101, p<0.0001), as illustrated in panel C of Figure 8. This is consistent with our prediction that a heavy final syllable would have increased amplitude because weight is a word-level property and therefore amplitude is available as a cue.

In summation, we find that closed final syllables, like in /navɛt/ *navette* 'shuttle', typically have higher pitch, amplitude and duration than open final syllables, like in /navɛ/ *navet* 'turnip',

consistent with Prediction 2. In the next section, we examine the results for vowel weight to determine whether weight effects are also found for vowels.

Figure 8: Results for final coda weight.

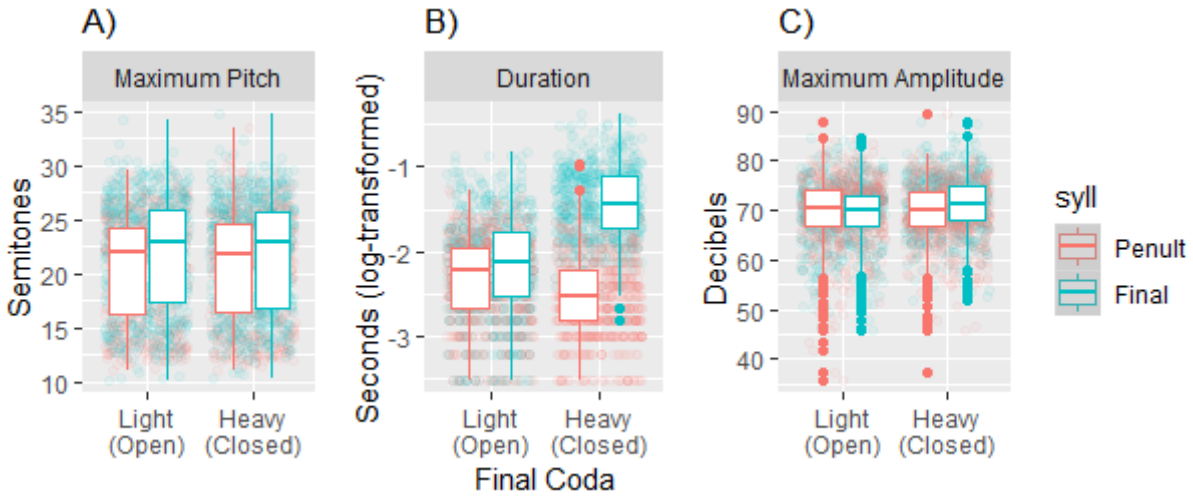

2.4.3 Vowel Weight

Prediction 3 stated that heavy vowels should attract prominence in penults, but that final open syllables should pattern as light and therefore vowel weight should not have a significant effect in this position. We predicted that RVs would be higher for heavy penults, but unaffected by heavy final-syllable vowels unless the final syllable is also closed, in which case the vowel would be more likely to attract prominence. Since we are discussing main effects here, we predicted that heavy penult vowels would attract prominence (leading to positive RVs), but we did not predict that heavy final-syllable vowels would be significantly different from light ones. We again begin by presenting the results for penults first (section 2.4.3.1) and then turn to final syllables (section 2.4.3.2).

2.4.3.1 Penult Vowel Weight

We predicted that heavy final syllables would be associated with higher RVs for pitch maxima, but the results are marginal (β=0.2493, p=0.0720). We illustrate these cues in panel A of Figure

9. Once again, asymmetries in the morphological and phonological profiles of the words in the data largely obscure the statistical trends visually.

We find the predicted increase in duration RVs when the penult vowel is heavy (β=0.1881, p=0.0024), shown in panel B of Figure 9, but find no significant effect of penult vowel weight on amplitude RVs, which may reflect our expectation that a small difference in duration may be present without signifying a difference in weight. As shown in panel C of Figure 9, the amplitude differences are small with a possible trade-off between syllables and therefore this statistical trend should be explored further in future work. In other words, a penultimate heavy vowel, as in /gɑto/ *gâteau* 'cake', is likely to have longer duration than the penultimate light vowel in a word like /bato/ *bateau* 'boat' and is likely to have a higher pitch maximum, but may not have higher amplitude, overall in line with Prediction 3.

Figure 9: Results for penult vowel weight.

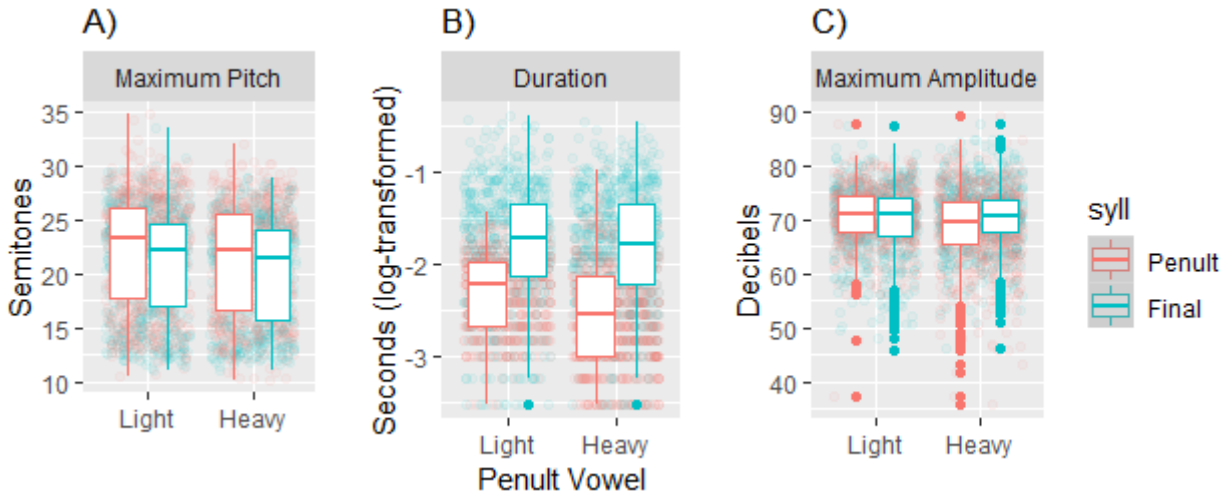

2.4.3.2 Final-Syllable Vowel Weight

For final syllable vowel weight, we predicted no main effects with the possible exception of a small increase in final rhyme duration. In this case, pitch maxima (panel A of Figure 10) and amplitude (panel C) show no noteworthy effect that does not result from other factors, primarily from the presence of a coda in the penult. Our models reveal no significant effects for these factors (though for the maximum pitch and amplitude there may be a small effect, which future work should test with a larger dataset). For duration, however, we do find that the RV is

significantly lower when the final syllable is heavy, consistent with the small predicted effect (β=-0.2662, p=0.0184). Overall, this means that the final syllables in words like /ami/ *ami* 'friend' (light final vowel) and /ane/ *année* 'year' (underlyingly heavy final vowel) may show no difference in pitch, duration or amplitude, consistent with Prediction 3.

Figure 10: Results for final vowel weight.

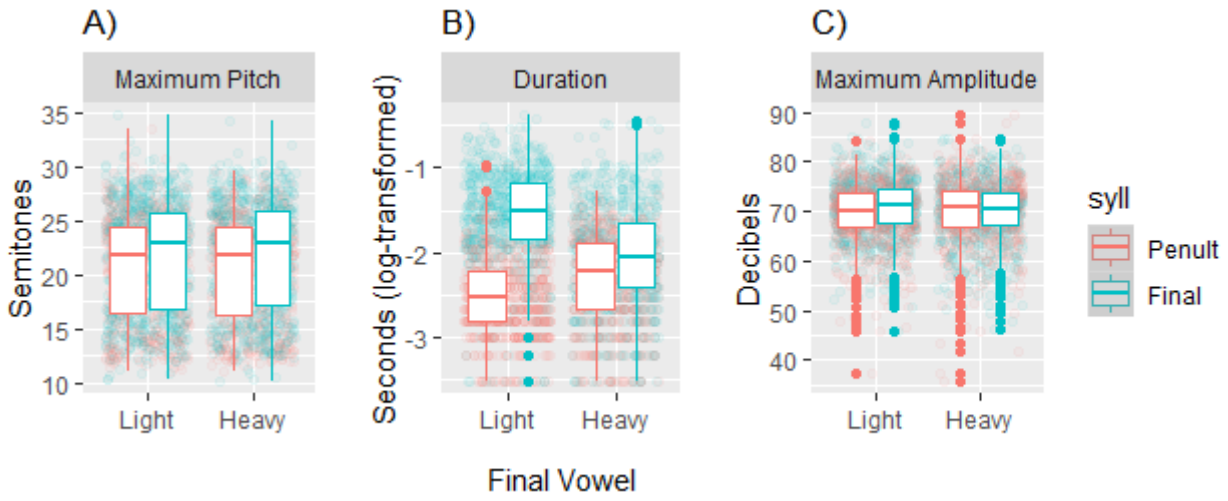

2.4.4 Superheavy Final Syllables

Prediction 3 stated that the final syllable's vowel weight effects would primarily be found in superheavy syllables because vowel length is retained in final closed syllables in French but not in final open syllables. We find no significant interaction, as shown in panel A of Figure 11. However, the figure suggests that future analyses should revisit this comparison, since closed final syllables with light vowels (leftmost in panel A) seem to show much smaller differences in cue values compared to those found in the other panels.

Turning to duration, shown in panel B of Figure 11, we find a significant interaction (β=0.4866, p=0.0006). However, while it seems that superheavy syllables are longer than regular closed syllables, the interaction predominantly reflects that the final syllable's weight is instead affecting duration in the penult.

As for amplitude RVs, we find a large and significant interaction (β=4.8909, p=0.0007) that, as shown in panel C of Figure 11, reveals that light syllables (i.e., those that are open and have a

light vowel) pattern differently than heavy syllables because the final syllable has lower amplitude than the penult when the final syllable has neither a coda nor a heavy vowel. In final open syllables that have a heavy vowel, the two final syllables have roughly equal amplitude on average, while closed final syllables on average have higher amplitude than the penult that precedes them.

Figure 11: Results for the interaction between vowel weight (x-axis) and coda weight (facets).

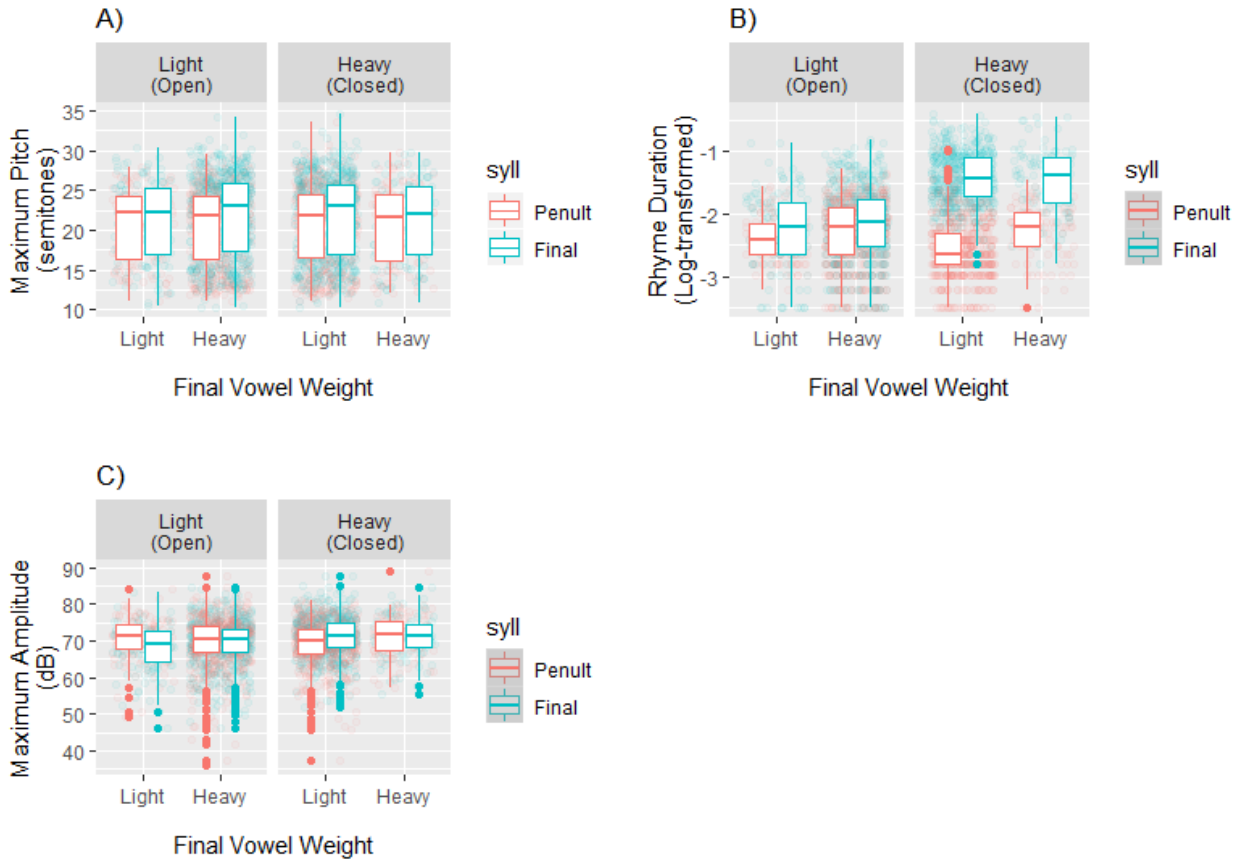

2.4.5 Summary of Results

In summation, we find that all acoustic cues are affected by the prosodic context and by weight. Regarding marking domains, our results are consistent with what was predicted, and therefore suggest that the basic patterns for marking prosodic domains in Laurentian French match the patterns for other dialects. Crucially, our predictions hold for weight: heavy syllables are associated with greater prominence than light syllables in all cases where weight is significant. In the following section, we discuss the implications of these results.

2.5 Discussion

Overall, we found the effects we predicted, with significant results both for prosodic domains and for weight, suggesting that both play a role in determining the realisation of prominence cues. We will begin the discussion of the results with the marking of prosodic domains in section 2.5.1. Following this, we will examine how prominence and its assignment signal weight in section 2.5.2. Finally, in section 2.5.3, we will discuss what these results suggest for classifying French typologically, as well as for the interpretation of results from previous studies.

2.5.1 Marking Prosodic Domains

In this section, we discuss what our results for prosodic domains suggest for marking domains, focusing both on the cues a speaker would intentionally manipulate and on possible perceptual cues (even if they are not intentionally manipulated for this purpose). We begin by discussing pitch in section 2.5.1.1, where we examine how the pitch maximum results reflect the pitch contours realised by speakers. In section 2.5.1.2, we turn to the relationship between lengthening and prosodic domains, focusing on phrase-final lengthening. Finally, in section 2.5.1.3, we discuss how our results for amplitude suggest that this acoustic cue may be used in perception, even though we do not find strong evidence of intentional manipulation of it to mark prosodic domains. Our results for pitch maxima largely coincide with both our predictions and the literature, but our results for duration and amplitude at first glance seem surprising. Overall, we will propose that the Laurentian French prosodic system for marking domains is like the one found in other dialects, as observed by Kaminskaïa (2009, 2015).

2.5.1.1 Pitch

We find that the IP's final syllable has significantly lower maximum pitch than the AP's final syllable does, which is consistent with Laurentian French speakers marking APs with a rising (LH*) bitonal unit and with IPs being assigned an additional low boundary tone (L%) that replaces any tone assigned to the AP-final syllable. This result suggests that, at least with respect to the general pitch contours, Laurentian French follows the same system as other dialects studied in the literature (with the exception of Parisian French where rising tones may be used to

mark IPs instead of only APs; Post and Delais-Roussarie, 2006). These results are also consistent with Prediction 1 that the right edges of prosodic domains are marked using pitch contours.

2.5.1.2 Duration

We found only a very small and barely significant rhyme duration difference between APs and IPs, with IPs tending to have longer final syllables compared to APs. That we did not find a robust result is consistent with the mixed results found in the literature for other dialects; the durational difference between APs and IPs may be very small, highly variable or non-existent, which leads to certain studies finding that final syllables get compressed in IPs, others finding that final syllables are further lengthened in IPs, and some studies not being able to conclude either way (cf. Demers, 2003; Jun and Fougeron, 2000; Michelas et al., 2010; Post and Delais-Roussarie, 2006; Simon, 2011). Given that the results across studies are so mixed, and based on our relatively marginal result, the prediction that higher domains will show greater degrees of lengthening cannot be confidently confirmed.

We conjecture, however, that the presence of mixed results across studies may be a consequence of the type of data analysed. If greater lengthening in IP-final tokens reflects planning limitations (with lengthening being a method to gain more time to plan upcoming words) or conversational cues (for example, signalling that the speaker is not ceding the floor), then perhaps it should be expected that IP-final lengthening will pattern differently across contexts. In the present study, which analysed read speech, the speakers do not need to plan the content of upcoming phrases (they need to retrieve lexical entries, but do not need to plan what the words themselves will be), meaning that there may be less need to slow down at the end of an IP to facilitate planning the next prosodic domain.

Future work on this corpus can test the proposal that lengthening is an effect of planning limitations, as the speakers provided both read speech and conversational speech. If this proposal is supported, it would suggest that the difference between domains is not a consequence of marking the right edges of domain types differently, and especially not of lengthening being proportional to domain size. Domain-final lengthening may nonetheless be used by listeners to chunk strings into prosodic phrases in regular conversation.

2.5.1.3 Amplitude

Lastly, we found that IPs have lower relative amplitude than APs do, seemingly contrary to our prediction that amplitude would not be used as a cue to phrasal prominence. While this may suggest that amplitude could be directly manipulated by speakers as a cue because a gradual decrease in amplitude could signal that the right edge of the current IP has not yet been reached, cross-linguistic evidence leads us to believe that amplitude is not intentionally used by speakers to mark the right edge of prosodic domains.

Based on findings from German (Poschmann and Wagner, 2015) and from Vietnamese (Brunelle, 2016) in addition to a similar proposal for French (see below), we suggest that the results obtained reflect aerodynamic and physiological effects. In particular, the articulatory force will be lowest IP-finally, leading to a decrease in amplitude unless the speaker intentionally counters these effects (e.g., to hold the floor). If the syllable that is assigned default prominence (signalled through higher pitch and longer duration) has lower amplitude and it is the result of a gradual decrease throughout the phrase, then it seems unlikely that amplitude is being intentionally manipulated to signal phrasal prominence, and that this reduction may instead be aerodynamic in nature.

Even if amplitude is not consciously manipulated by speakers, it may still be used as a perceptual cue by listeners. This proposal is not only consistent with the cross-linguistic acoustic work just mentioned, but we believe it is also supported by the results of a previous perceptual study on French speakers. Although Féry (2011) contends that amplitude is not a possible cue to prominence in French, and amplitude is typically not tested in perception studies of French prominence, Schwab and Llisterri (2013) found that French speakers learning Spanish readily attended to amplitude to identify stressed syllables in Spanish. It may be that the usefulness of amplitude for parsing phrases in French makes it a possible candidate for transfer: French speakers could repurpose a cue for identifying IP boundaries in French to locate stress in Spanish. However, as we will discuss in the next section, it may be that the relationship that amplitude has with weight – rather than its correlation with phrasal boundaries – is what makes it a good candidate for locating word-level stress in Spanish.

2.5.2 Signalling Weight

Our results provide support for the hypothesis in (1) that prominence assignment in French is sensitive to weight. Only two studies to our knowledge have quantitatively examined relationships between weight and prominence in French. The first (Paradis and Deshaies, 1990) is a perceptual study on Laurentian French that found that listeners were more likely to categorise a syllable as prominent if it is closed. The second (Thibault and Ouellet, 1996) demonstrates that the pitch contours that result from prominence shifting to the penult (using heavy vowels in their study to elicit these tokens) are distinct from those that arise under focus in Laurentian French, and therefore that penultimate prominence cannot be explained by focus. The current study, we believe, is the first to examine both vowel and coda weight when probing prominence assignment, as well as the first to test weight effects in speech directly. The results align with earlier work motivating the existence of weight contrasts based on segmental processes in French (e.g., Armstrong, 1999; Scullen, 1997; but cf. Bullock, 1994).

Beginning with coda weight, which was expected to significantly attract prominence based on our second prediction, we observe that the presence of a final coda increases the relative prominence of the final syllable, affecting pitch, amplitude, and duration. Similarly, closed penults show an increase in relative value for these same cues. Our results suggest not only that these cues signal weight,[22] but additionally, that only one syllable is targeted by these effects, and the other may even show *decreased* values for the cues. These results are consistent with a phonological representation of weight, like the mora (Hayes, 1989; Hyman, 1985); relative durations are computed by comparing the weights of the final two syllables. In Section 2.2.1, we noted that there is debate in the literature on French about whether word-final consonants are truly codas or whether they are onsets of empty-headed syllables. Given that word-final consonants in the data we examined bear weight, we conjecture that they are best analysed as codas, but we leave further testing of alternative analyses (e.g., that some word-final consonants may pattern as onsets) for future work. We also leave for future work the question of whether lengthening processes (like those discussed in footnote 19) and the shortening of underlyingly heavy vowels in final open syllables reflect adding and losing morae, respectively.

[22] With the exception of cases where the intonational contour requires a different tone target, like in the case of heavy IP-final syllables, which have low pitch because of the IP's boundary tone.

There are trends for heavy penult vowels to attract prominence, but underlying vowel weight in the final syllable is not sufficient to preserve prominence on that syllable. These results suggest that vowel weight (when not in a final open syllable) contributes to syllable weight. Final heavy vowels slightly enhance the effects of being closed – that is, either making that syllable even more prominent or further decreasing the likelihood that prominence shifts to the penult, which suggests that those syllables may be phonologically heavier than other closed syllables and, thus, that the label superheavy may be appropriate. This result is particularly noteworthy because it confirms that underlyingly heavy vowels in final open syllables pattern as short for prominence assignment, which accords with the inability to diphthongise oral vowels in those syllables; only heavy oral vowels in final closed syllables lengthen and can diphthongise.

Returning to Table 4 where we had classified vowels as light or heavy, we included schwa as a light vowel. We did not analyse rhymes with schwa, however, as none were present in the positions targeted in the data. In the literature, schwa is not assigned prominence (Garde, 1968; Pasdeloup, 1990; Prieto et al., 2005), rather than simply being disfavoured in prominence assignment like the light vowels in our data. This is consistent with schwa having no mora (e.g., Hyman, 1985; Tranel, 1984; but cf. Eychennes, 2006). Turning to the heavy vowels in Table 4, we included nasal vowels in the set of heavy vowels, though nasal vowels differ from oral ones because they appear to preserve their long duration and can undergo diphthongisation in final open syllables (see footnote 7). Since historically nasal vowels were derived from a nasal consonant, we leave for future work whether syllables containing nasal vowels conserve their weight like closed syllables.

In summation, we have found evidence of weight effects for vowels as well as codas for all three acoustic cues, consistent with our predictions, but we observe that these effects are not identical. Heavy vowels only pattern as heavy (i.e., attract prominence) when they are not word-final, while codas show the same prominence-attracting property in both penultimate and final syllables. As such, while prominence assignment is probabilistic, it is not arbitrary when prominence shift is most likely to occur.

2.5.3 Implications for the Prominence System

Our results demonstrate that the marking of prosodic domains matches what has been found for other dialects of French. The cues used to mark prosodic domains also signal weight, which means that these factors interact to produce the prominence patterns we observe in French, or at least in Laurentian French.[23] Heavy syllables attracting prominence therefore has important repercussions for our understanding of the prominence system itself in (this variety of) French.

At the very least, prominence – including the assignment of the AP's H* tone – appears to play a different role in the grammar of French than conventionally proposed; *word*-level factors (i.e., weight) influence the prosody of a *phrase*, while previously it was generally assumed that only phrase-level information was relevant (barring the inability to assign prominence to a phrase-final schwa). This may account for speakers' judgments in perceptual studies being variable and difficult to interpret (e.g., Frost, 2011; Paradis and Deshaies, 1990) and may have led to the characterisation of French speakers as "stress deaf" (e.g., Dupoux et al., 1997; Peperkamp and Dupoux, 2002). If only the location of phrasal domain edges is presumed to be relevant for prominence assignment and stimuli are resynthesised without taking word-level factors into account, then we might expect speakers, faced with conflicting acoustic information and uncertainty as to which aspect of prominence they are asked to identify, to provide inconsistent responses in experimental settings.

The role of word-level factors may also explain previous findings like those of Frost (2011), which showed that French speakers' perception of prominence was affected not only by the pitch contour of the stimuli, but also by the specific word – a result that we expect when taking into account that the stimuli differed in vowel weight. Additionally, the fact that amplitude is manipulated as a cue – even though this is for signalling weight rather than for marking domains – could also help explain why French speakers listening to Spanish readily use amplitude as a cue to locate Spanish stress (Schwab and Llisterri, 2012).

Signalling word-level factors (weight) using the same cues as those used to mark prosodic domains has greater implications than simply explaining otherwise surprising results in perceptual studies. In particular, it helps shed light on the type of prominence system that French

[23] The current study does not have sufficient speakers to confidently test inter-speaker variability, but the consistency within our speaker sample suggests the weight effects are found across speakers.

employs. In section 2.2.3, we mentioned that APs are marked with an LH* tone, described as a pitch accent, but we did not elaborate on the matter. In the literature, there is little discussion about why it is referred to as a pitch accent: it is assumed to fall on the last non-schwa vowel in the AP like a boundary tone. It may be that the pitch accent notation is used because the tone is assigned by the domain that cross-linguistically assigns pitch accents (as per Gordon, 2014). In other words, the pitch accent notation in French reflects the *domain* of assignment but may not be meant to indicate that the tone is *formally* a pitch accent: pitch is used to cue a single lexically prominent syllable, with the phonological assignment occurring in the phrasal domain.

As previously noted, the formal description of obligatory final prominence in French is debated, and often is not clarified. On one hand, many studies refer to the prominence in French as being stress (e.g., Cutler, 2005; Schwab and Llisterri, 2012), which then should be assigned at the level of the word (as has been proposed for Midi French; Sichel-Bazin, 2016) and therefore its location would be expected to be sensitive to word-level properties, notably weight. On the other hand, French prominence is often described as phrasal or post-lexical, and thus not sensitive to word-level properties (e.g., Féry, 2013; Jun and Fougeron, 1995). In this paper, however, we have shown that the H* tone patterns like a pitch accent: the H* in the AP's LH* tone is attracted to the rightmost heavy syllable within a limited window, which leads to the pitch peak being on heavy penults when the final syllable is light.[24] This is consistent with what we observe in pitch accent systems cross-linguistically (Gordon, 2014). We therefore propose based on this and on the acoustic cues manipulated (notably amplitude) that French may be best categorised as a pitch accent language.

2.6 Conclusion

The results of this study confirm that prominence shifts do occur in French, with the acoustic cues associated with prominence being realised on the penult. In this analysis, we show that, although this phenomenon is probabilistic such that the location of prominence cannot

[24] As Jun and Fougeron (1995) found that a small proportion of their tokens showed antepenultimate prominence, we expect that a window of two or three syllables from the right edge is considered when assigning prominence. We leave this question for future research, since to our knowledge the contexts allowing antepenultimate prominence are not known and it may be that another factor – like morphological structure – allows the prominence to surface earlier in the AP-final word than might otherwise be expected.

consistently be predicted for any given token, it is not arbitrary once we examine the broader patterns. In particular, weight effects predict the prominence assignment patterns we observe and may also explain the prominence shift noted in previous studies that examined data from as long as a century ago (Martin, 2011).

One question of considerable debate has been how to formally categorise the prosodic system in French. The system may itself be in flux (as suggested by Fónagy, 1980), leading to the variation in prosodic realisations found within and across dialects, and we aim to determine what structure underlies or constrains this variation. We expect that weight effects are not confined to Laurentian varieties of French, but instead that weight sensitivity contributes to prominence shifts across varieties and should be robustly found across speakers. It could be, for example, that the rate of prominence shift is markedly different in dialects like Parisian French where many vowel contrasts have been or are being lost (e.g., Berit Hansen, 2012). Given that different dialects have different phonological contrasts, future work on other varieties may additionally (a) test for weight sensitivity, (b) compare dialects' vowel inventories with their prominence shift patterns, and (c) verify whether (word-final) codas contribute to weight.

We may infer that there is a trade-off whereby word-level factors are signalled at the expense of marking prosodic domains, but we expect that this is not the case. The word containing the phrasal prominence is still on the right edge of the domain and locating the pitch accent near this edge allows listeners to interpret both the word's phonological profile and the word's position in the phrase. Signalling word-level prominence may therefore not interfere with marking phrase-level prominence because the availability of prominence shift is largely predictable to listeners. Listeners should still be able to parse the word and phrase boundaries successfully because they – like speakers – can compute whether prominence shift could occur.

References

Armstrong, S. D. (1999). *Stress and Weight in Québec French*. Master's thesis: University of Calgary.

Avanzi, M., A. Lacheret, and B. Victorri. (2010). A Corpus-Based Learning Method for Prominences Detection in Spontaneous Speech. *Proceedings of Prosodic Prominence, Speech Prosody 2010 Satellite Workshop*, 1-4.

Avanzi, M., N. Obin, A. Bardiaux, and G. Bordal. (2011). *Données et hypothèses sur la variation prosodique de 6 variétés de français parlées en France*. Journées PFC, Paris.

Avanzi, M., S. Schwab, J.-P. Goldman, P. Montchaud, I. Racine, and H. Andreassen. (2011). *Étude acoustique de l'accentuation pénultième dans trois variétés de français*. Journées PFC, Paris.

Baraduc, J., G. Bergounioux, V. Castellotti, C. Dumont, and M.-H. Lansari. (1989). Le statut linguistique des voyelles moyennes. *Langage et société,* 49, 5-24.

Bardiaux, A., and P. Mertens. (2014). Normalisation des contours intonatifs et étude de la variation régionale en français. *Nouveaux cahiers de linguistique française,* 31, 273-284.

Bates, D., M. Maechler, B. Bolker, and S. Walker. (2015). Fitting Linear Mixed-Effects Models in lme4. *Journal of Statistical Software,* 67(1), 1-48.

Benua, L. (1997). *Transderivational Identity: Phonological Relations between Words*. Doctoral thesis: University of Massachusetts Amherst.

Berit Hansen, A. (2012). A Study of Young Parisian Speech: Some Trends in Pronunciation. In Gess, R., C. Lyche, and T. Meisenburg (Eds.), *Phonological Variation in French: Illustrations from Three Continents* (151-172). Amsterdam: John Benjamins.

Beyssade, C., E. Delais-Roussarie, and J.-M. Marandin. (2007). The prosody of interrogatives in French. *Cahiers de Linguistique Française* 28: 163-175.

Boersma, P., and D. Weenink. (2015). *Praat: Doing Phonetics by Computer*. [Computer program]. Retrieved Sept. 24, 2015, http://www.praat.org/.

Brunelle, M. (2016). *When Intonation Fails to Phonologize: The Case of Southern Vietnamese*. Montreal-Ottawa-Laval-Toronto Phonology Workshop, Carleton University (Canada).

Bullock, B. E. (1994). Does the French syllable have weight? In: M. Mazzola, (ed.), *Issues and Theory in Romance Linguistics*, Georgetown: Georgetown University Press, 3-19.

Carton, F., M. Rossi, D. Autesserre, and P. Léon. (1983). *Les accents du français*. Paris: Hachette.

Centre national de ressources textuelles et lexicales (CNRTL). (2012). Lancement. In *Portail lexical: étymologie*. Online. Consulted March 3rd 2016. http://www.cnrtl.fr/etymologie/lancement

Chitoran, I. (1996). Prominence vs. Rhythm: The Predictability of Stress in Romanian. In Zagona, K. (Ed.), *Grammatical Theory and Romance Languages: Selected Papers from the 25th Symposium on Romance Languages* (47-58). Amsterdam: John Benjamins.

Coquillon, A., and G. Turcsan. (2012). An Overview of the Phonological and Phonetic Properties of Southern French. In Guess, R., C. Lyche, and T. Meisenburg (Eds.), *Phonological Variation in French: Illustrations from Three Continents* (105-128). Amsterdam: John Benjamins.

Coquillon, A. (2005). *Caractérisation prosodique du français de Marseille*. Doctoral thesis: Université Aix-Marseille I.

Côté, M.-H. (2014, April). *Phonologie du français contemporain: un outil pour l'analyse et l'enseignement du français*. Paper presented to the Department of Languages, Linguistics and Translation, Laval University, Quebec (Canada).

Côté, M.-H. (2012). Laurentian French (Québec): Extra Vowels, Missing Schwas and Surprising Liaison Consonants. In Gess, R., C. Lyche, and T. Meisenburg (Eds.), *Phonological Variation in French: Illustrations from Three Continents* (235-274). Amsterdam: John Benjamins.

Côté, M.-H. (2011). Final Consonants. In M. van Oostendorp, C. J. Ewen, E. Hume, and K. Rice, *The Blackwell Companion to Phonology: Volume I* (848-872). Oxford: Blackwell Publishing.

Cutler, A. (2005). Lexical Stress. In P. Pisoni, and R. Remez, *The Handbook of Speech Perception* (264-289). Oxford: Blackwell Publishing.

Delais-Roussarie, E. (1996). Phonological Phrasing and Accentuation in French. In : M. Nespor, and N. Smith, (eds.), *Dam Phonology : HIL Phonology Paper II*, Holland: Holland Academic Graphics, 1-38.

Delattre, P. (1939). Accent de mot et accent de groupe. *The French Review,* 13(2), 141-146.

Delattre, P. (1968). L'accent final en français: accent d'intensité, accent de hauteur, accent de durée. *The French Review*, 12, 141-145.

Dell, F. (1995). Consonant Clusters and Phonological Syllables in French. *Lingua*, 95, 5-26.

Demers, M. (2003). La voix du plus fort. Étude acoustique sur le registre vocal en tant qu'indicateur sociolectal et dialectal en français spontané. In Demers, M. (Ed.), *Registre et voix sociale*, 196.

Di Cristo, A. (1998). Intonation in French. In Hirst, D., and A. Di Cristo (Eds.), *Intonation Systems: A Survey of Twenty Languages* (195-218). Cambridge: Cambridge University Press.

Di Cristo, A., and D. Hirst. (1993a). Rythme syllabique, rythme mélodique et représentation hierarchique de la prosodie du français. *Travaux de l'Institut de Phonétique d'Aix,* 15, 9-24.

Di Cristo, A., and D. Hirst. (1993b). Prosodic Regularities in the Surface Structure of French Questions. *Proceedings of the ESCA Workshop on Prosody (Lund, Sweden)*, 268-271.

Di Cristo, A., and D. Hirst. (1996). Vers une typologie des unités intonatives du français. *XXIème Journées Européennes du Patrimoine (Avignon, France),* 219-222.

Drager, K., and J. Hay. (2012). Exploiting random intercepts: Two case studies in sociophonetics. *Language Variation and Change*, 24(1), 59-78.

Dumas, Denis. (1974). Durée vocalique et diphtongaison en français québécois. *Cahier de linguistique,* 4, 13-55.

Dupoux, E., C. Pallier, N. Sebastian, and J. Mehler. (1997). A Destressing 'Deafness' in French? *Journal of Memory and Language,* 36, 406-421.

Dupoux, E., S. Peperkamp, and N. Sebastian-Galles. (2008). Limits on Bilingualism Revisited: Stress 'Deafness' in Simultaneous French-Spanish Bilinguals. *Cognition*, 106(2), 682-706.

Dupoux, E., N. Sebastian-Galles, E. Navarete, and S. Peperkamp. (2007). Persistent Stress 'Deafness': The Case of French Learners of Spanish. *Cognition*, 106(2), 682-706.

Durand, J., B. Laks, and C. Lyche. (2002). La phonologie du français contemporain: usages, variétés et structure. In C. Pusch, and W. Raible (Eds.), *Romanistische Korpuslinguistik - Korpora und Gesprochene Sprache* (93-106). Tübingen: Gunter Narr Verlag.

Durand, J., B. Laks, and C. Lyche. (2009). Le Projet PFC: une source de données primaires structurées. In J. Durand, B. Laks, and C. Lyche (Ed.), *Phonologie, variation et accents du français* (19-61). Paris: Hermès.

Eychenne, J. (2006). *Aspects de la phonologie du schwa dans le français contemporain. Optimalité, visibilité prosodique, gradience*. Doctoral thesis: Université de Toulouse-Le Mirail.

Féry, C. (2013). Focus as Prosodic Alignment. *Natural Language, and Linguistic Theory,* 31, 683-734.

Féry, C., R. Hörnig, and S. Pahaut. (2010). Phrasing in French and German: an Experiment with Semi-Spontaneous Speech. In Gabriel, C., and C. Lleó (Eds.), *Intonational Phrasing at the Interfaces: Cross-Linguistic and Bilingual Studies in Romance and Germanic* (11-41). Amsterdam: John Benjamins.

Fónagy, J. (1983). *La Vive voix: essais de psycho-acoustique*. Paris: Payot.

Fónagy, I. (1980). L'accent français, accent probabilitaire: dynamique d''un changement prosodique. I. Fónagy, and P. Léon (Eds.), *L'accent en français contemporain, Studia Phonetica,* 15, 123-233.

Fónagy, I., (1979). L'accent français: accent probabilitaire. In I. Fónagy, and P. R. Léon (Eds.), *L'accent en français contemporain* (123-133). Montréal/Paris/Bruxelles: Didier.

Fónagy, J. (1976). Prosodie professionnelle et changements prosodiques. *Le français moderne XLIV-3* (193-228). Paris: Editions d'Artrey.

Frost, D. (2011). Stress and Cues to Relative Prominence in English and French: A Perceptual Study. *Journal of the International Phonetic Association,* 41(1), 67-84.

Garcia, G. (2017). Weight Gradience and Stress in Portuguese. *Phonology,* 34(1), 41-79.

Garde, P. (1968). *L'accent*. Paris: Presses Universitaires de France.

Gendrot, C., Gerdes, K., and Adda-Decker, M. (2011). Impact of Prosodic Position on Vocalic Space in German and French. In *Proceedings of the 17th International Congress of Phonetic Sciences*, 731-734.

Georgeton, L., and J. Fougeron. (2014). Domain-Initial Strengthening on French Vowels and Phonological Contrasts: Evidence from Lip Articulation and Spectral Variation. *Journal of Phonetics,* 44 (May), 83-95.

Goad, H., and A.-E. Prévost. (2011). *A Test Case for Markedness: The Acquisition of Québec French*. Ms., McGill University.

Goldman, J.-P., and A.-C. Simon. (2007). *La variation prosodique régionale en français (Liège, Vaud, Tournai, Lyon)*. *Regards* croisés sur la phonologie du français contemporain*,* Paris.

Gordon, M. (2005). A Perceptually-Driven Account of Onset-Sensitive Stress. *Natural Language, and Linguistic Theory,* 23, 595-653.

Gordon, M. (2014). Disentangling Stress and Pitch Accent: Toward a Typology of Prominence at Different Prosodic Levels. In H. van der Hulst (Ed.), *Word Stress: Theoretical and Typological Issues* (83-118), Oxford: Oxford University Press.

Grammont, M. (1914). *Traité de phonétique*. Paris: Delagrave.

Halle, M. (1973). Stress Rules in English: A New Version. *Linguistic Inquiry,* 4, 451-464.

Hannahs, S. J. (1995). *Prosodic Structure and French Morphophonology*. Tübingen: Linguistische Arbeiten.

Hayes, B. (1982). Extrametricality and English Stress. *Linguistic Inquiry,* 13(2), 227-276.

Hayes, B. (1989). Compensatory Lengthening in Moraic Phonology. *Linguistic Inquiry,* 20, 253-306.

Hayes, Bruce. (1995). *Metrical Stress Theory: Principles and Case Studies*. Chicago: Chicago University Press.

Hirst, D., and A. Di Cristo. (1984). French Intonation: a Parametric Approach. *Die Neueren Sprache,* 83(5), 554-569.

Hirst, D., and A. Di Cristo. (1996). Y a-t-il des unités tonales en français? *XXIème Journées Européennes du Patrimoine (Avignon, France),* 223-226.

Hirst, D., A. Di Cristo, and R. Espesser. (1998). Levels of Description and Levels of Representation in the Analysis of Intonation. In D. Hirst, and A. Di Cristo (eds). *Intonation Systems: A Survey of Twenty Languages* (51-87). Dordrecht: Kluwer Academic Publisher.

Hyman, L. (1985). *A Theory of Phonological Weight*. Dordrecht: Foris.

Jun, S-A, and C. Fougeron. (2002). Realizations of Accentual Phrase in French Intonation. *Probus,* 14, 147-172.

Jun, S.-A., and C. Fougeron. (2000). A Phonological Model of French Intonation. In A. Botinis (ed.), *Intonation: Analysis, Modeling and Technology* (209-242). Dordrecht: Kluwer Academic Publisher.

Jun, S-A, and C. Fougeron. (1995). The Accentual Phrase and the Prosodic Structure of French. *Proceedings of the International Congress of Phonetic Science (Stockholm, Sweden),* 2, 722-725.

Kaminskaïa, S. (2015). Tonal Patterns, Associations, and Alignment of Peaks in Regional French. *Word*, 61(2), 101-140.

Kaminskaïa, S. (2009). *La variation intonative dialectale en français: une approche phonologique*. Münich: Lincom Europa Verlag.

Lamontagne, J. (2014). *The realisation of mid vowels in Laurentian French: Processes affecting the final syllables*. MA memoire, University of Ottawa.

Lamontagne, J., H. Goad, and M. Sonderegger. (2017). *The continuation contour in French: Realisation and representation*. 2nd Intonation Workshop in Toronto (Canada).

Légaré, L. (1978). Relâchement des voyelles hautes et reformulation. *Cahier de linguistique*, 7, 31-42.

Léon, P., and P. Bhatt. (2005). *Structure du français moderne*. Toronto, Canada: Canadian Scholars' Press.

Liberman, M., and A. Prince. (1977). On Stress and Linguistic Rhythm. *Linguistic Inquiry,* 8, 249-336.

Martin, P. (2011). *La structure prosodique cognitive: réalisations régionales*. Journées PFC, Paris.

Martin, P. (2004). L'intonation de la phrase dans les langues romanes: l'exception du français. *Langue française,* March 2004, 36-55.

Martin, P. (1987). Prosodic and Rhythmic Structures in French. *Linguistics,* 25(5), 925-949.

Martin, P. (1975). Analyse phonologique de la phrase française. *Linguistics,* 146, 35-68.

Mertens, P. (1993). Intonational Grouping, Boundaries, and Syntactic Structure in French. In House, D., and P. Touati (Eds.), *Proceedings of the ESCA Workshop on Phonology (Sept. 27-29, 1993, Lund),* 156-159.

Mertens, P. (2006). A Predictive Approach to the Analysis of Intonation in Discourse in French. In Kawaguchi, Y., Fonagy, I. and T. Moriguchi, T., (Eds.), *Prosody and Syntax. Usage-Based Linguistic Informatics* 3 (64-101). Amsterdam: John Benjamins.

Mertens, P. (1987). *L'Intonation du français*. PhD Dissertation, Katholieke Universiteit Leuven.

Michelas, A., and M. D'Imperio. (2011). Uncovering the Role of the Intermediate Phrase in the Syntactic Parsing of French. *Proceedings of the International Congress of Phonetic Science in Hong Kong, China (17-21 August, 2011),* 1374-1377.

Milne, P. (2014). *The Variable Pronunciations of Word-Final Consonant Clusters in a Force Aligned Corpus of Spoken French*. Doctoral thesis: University of Ottawa.

Montreuil, J.-P. (1995). Weight and Length in Conservative Regional French. *Lingua,* 95, 77–96.

Morin, Y.-C. (2002). Les premiers immigrants et la pronunciation du français au Québec. *Revue québécoise de linguistique,* 31(1), 39-78.

Ouellet, M., and B. Tadif. (1996). From Segmental Duration Properties to Rhythmic Structure: A Study of Interactions between High and Low Level Constraints. In *Proceedings of the 4th International Conference on Second Language Processing*, 1177-1180.

Paradis, C. (1985). *An Acoustic Study of Variation and Change in the Vowel System of Chicoutimi and Jonquière (Québec).* Doctoral thesis: University of Pennsylvania.

Pasdeloup, V. (1990). *Modèle de règles rythmiques du français appliquées à la synthèse de la parole*. Doctoral thesis: Université de Provence.

Peperkamp, S., and E. Dupoux. (2002). A Typological Study of Stress 'Deafness'. In Gussenhoven, C., and N. Warner (Eds.), *Laboratory Phonology,* 7 (203-240). Berlin: Mouton de Gruyter.

Peperkamp, S., I. Vendelin, and E. Dupoux. (2010). Perception of Predictable Stress: A Cross-Linguistic Investigation. *Journal of Phonetics,* 38(3), 422-430.

Pilch, H. (1972). La théorie de la phonologie. In *Proceedings of the Fifth International Congress of Phonetic Sciences (Montréal 1971)* (105-119). Basel: Karger.

Portes, C., and R. Bertrand. (2005). Some Cues about the Interactional Value of the 'Continuation' Contour in French. In *Proceedings of Discourse-Prosody Interface Symposium - IDP05* (Seridisc, Belgium), 1-14.

Post, B. (2000). *Tonal and Phrasal Structures in French Intonation*. The Hague: Holland Academic Graphics.

Post, B. (1993). *A Phonological Analysis of French Intonation*. MA thesis, University of Nijmegen.

Post, B., E. Delais-Roussarie, and A.C Simon. (2006). Développer un système de transcription des phénomènes prosodiques. In Simon, A. C., and G. Caelen (Eds.), *Bulletin PFC*, 51-68.

Prieto, P. (2005). Syntactic and Eurhythmic Constraints on Phrasing Decisions. In M. Horne, and M. van Oostendorp (Eds.), *Studia Linguistica* 59: *Boundaries in Intonational Phonology (*194-222). Oxford: Blackwell Publishing.

Prieto, P. (1997). Prosodic Manifestation of Syntactic Structure in Catalan. In Martínez-Gil, F., and Morales-Front, A. (Eds.), *Issues in the Phonology of the Iberian Languages* (179-199). Washington, D.C.: Georgetown University Press.

Prieto, P., M. D'Imperio, and B. Gili Fivela. (2005). Pitch Accent Alignment in Romance: Primary and Secondary Associations with Metrical Structure. *Language and Speech,* 48(4), 359-396.

Prince, A. (1990). Quantitative Consequences of Rhythmic Organization. *Chicago Linguistic Society,* 26(2), 355-398.

R Development Core Team. (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, http://www.R-project.org.

Rossi, M. (1980). Le français, langue sans accent? *Studia Phonetica,* 15, 85-101.

Ryan, K. (2016). Phonological Weight. *Language and Linguistics Compass,* 10, 720-733.

Santiago, F. (2011). *Les relations entre phrasing et contours intonatifs en français*. Journées PFC, Paris.

Schwab, S., and J. Llisterri. (2012). The Role of Acoustic Correlates of Stress in the Perception of Spanish Accentual Contrasts by French Speakers. In Ma, Q., H. Ding, and D. Hirst (Eds.), *Proceedings of the 6th International Conference on Speech Prosody (Shanghai, May 22-25, 2012),* 1, 350-353.

Scullen, M.-E. (1997). *French Prosodic Morphology: A Unified Account*. Bloomington: Indiana University Linguistics Club.

Selkirk, E. (1972). *The Phrase Phonology of English and French*. New York: Garland Publishing Co.

Shattuck-Hufnagel, S., and A. Turk. (1998). The Domain of Phrase-Final Lengthening in English. *The Journal of the Acoustical Society of America,* 103(5), 2889.

Sichel-Bazin, R. (2016). Le français du Midi: une variété de français à accent lexical. *Langue française,* 191(3), 47-66.

Sichel-Bazin, R., C. Buthke, and T. Meisenburg. (2011). La *prosodie du français parlé à Lacaune (Tarn): influences du substrat occitan*. Journées PFC, Paris.

Simon, C. (2011). La prosodie des accents régionaux en français. État des lieux. *Journées PFC,* Paris.

Simon, A-C. (2004). *La structuration prosodique du discours en français*. Bern: Peter Lang.

Skoruppa, K., F. Pons, A. Christophe, L. Bosch, E. Dupoux, N. Sebastián-Gallés, R. A. Limissuri, and S. Peperkamp. (2009). Language-Specific Stress Perception by Nine-Month-Old French and Spanish Infants. *Developmental Science,* 12(6), 914-919.

Smith, C. (2012). Emphasis does Not Always Coincide with Phrasal Boundaries in Spontaneous Spoken French. *Proceedings of JEP-TALN-RECITAL in Grenoble (June 4-8, 2012),* 1, 17–24.

Statistics Canada. (2012). Focus on Geography Series, 2011 Census. *Statistics Canada Catalogue* no. 98-310-XWE2011004. Ottawa, Ontario. Analytical products, 2011 Census. Last updated October 24, 2012.

Thibault, L., and M. Ouellet. (1996). Tonal Distinctions between Emphatic Stress and Pretonic Lengthening in Quebec French. *Proceedings of the Fourth International Conference on Spoken Language,* 2, 638-641.

Togeby, K. (1965). *Structure immanente de la langue française*. Paris: Larousse.

Tranel, B. (1984). Closed Syllable Adjustment and the Representation of Schwa in French. *Proceedings of the Berkeley Linguistics Society 10* (65-75). Berkeley: Berkeley Linguistics Society.

Vaissière, J. (2010). Le français, langue à frontières par excellence. In Delomiier, D., and Morel, M-A. (Eds.), *Frontières, du linguistique au sémiotique*, 10-20.

Vaissière, J. (1991). Rhythm, Accentuation and Final Lengthening. In Sundberg, J., L. Nord, and R. Carlson (Eds.), *French in Music, Language, Speech and Brain,* Wenner-Gren International Symposium Series Macmillan Press, vol. 59, 108-120.

Vaissière, J. (1974). On French Prosody. *Research Laboratory of Electronics Progress Report*, Massachusetts Institute of Technology, 115, 212-223.

Vaissière, J., and A. Michaud. (2006). Prosodic Constituents in French: A Data-Driven Approach. In Fónagy, I., Y. Kawaguchi, and T. Moriguchi (Eds.), *Prosody and Syntax* (47-64). Amsterdam: John Benjamins.

Walker, D. C. (1984). *The Pronunciation of Canadian French*. Ottawa, Canada: University of Ottawa Press.

Welby, P. (2007). The Role of Early Fundamental Frequency Rises and Elbows in French Word Segmentation. *Speech Communication,* 49(1), 28-48.

Welby, P. (2006). French Intonational Structure: Evidence from Tonal Alignment. *Journal of Phonetics,* 34(3), 343-371.

Welby, P. (2004). The Structure of French Intonational Rises: A Study of Text-to-Tune Alignment. In B. Bel, and I. Marlien (Eds.), *Proceedings of the Speech Prosody 2004 Conference*, 127-130.

Zec, D. (1995). Sonority Constraints on Syllable Structure. *Phonology,* 12, 85-129.

Appendix A

Read text in the original French:

Le village de Beaulieu est en grand émoi. Le Premier Ministre a en effet décidé de faire étape dans cette commune au cours de sa tournée de la région en fin d'année. Jusqu'ici les seuls titres de gloire de Beaulieu étaient son vin blanc sec, ses chemises en soie, un champion local de course à pied (Louis Garret), quatrième aux jeux olympiques de Berlin en 1936, et plus récemment, son usine de pâtes italiennes. Qu'est-ce qui a donc valu à Beaulieu ce grand honneur? Le hasard, tout bêtement, car le Premier Ministre, lassé des circuits habituels qui tournaient toujours autour des mêmes villes, veut découvrir ce qu'il appelle « la campagne profonde ».

Le maire de Beaulieu - Marc Blanc - est en revanche très inquiet. La cote du Premier Ministre ne cesse de baisser depuis les élections. Comment, en plus, éviter les manifestations qui ont eu tendance à se multiplier lors des visites officielles ? La côte escarpée du Mont Saint-Pierre qui mène au village connaît des barrages chaque fois que les opposants de tous les bords manifestent leur colère. D'un autre côté, à chaque voyage du Premier Ministre, le gouvernement prend contact avec la préfecture la plus proche et s'assure que tout est fait pour le protéger. Or, un gros détachement de police, comme on en a vu à Jonquière, et des vérifications d'identité risquent de provoquer une explosion. Un jeune membre de l'opposition aurait déclaré: « Dans le coin, on est jaloux de notre liberté. S'il faut montrer patte blanche pour circuler, nous ne répondons pas de la réaction des gens du pays. Nous avons le soutien du village entier. » De plus, quelques articles parus dans La Dépêche du Centre, L'Express, Ouest Liberté et Le Nouvel Observateur indiqueraient que des activistes des communes voisines préparent une journée chaude au Premier Ministre. Quelques fanatiques auraient même entamé un jeûne prolongé dans l'église de Saint Martinville.

Le sympathique maire de Beaulieu ne sait plus à quel saint se vouer. Il a le sentiment de se trouver dans une impasse stupide. Il s'est, en désespoir de cause, décidé à écrire au Premier Ministre pour vérifier si son village était vraiment une étape nécessaire dans la tournée prévue. Beaulieu préfère être inconnue et tranquille plutôt que de se trouver au centre d'une bataille politique dont, par la télévision, seraient témoins des millions d'électeurs.

English translation:

The village of Beaulieu is full of commotion. The Prime Minister has indeed decided to stop in this community during his tour of the region at the end of the year. Until now, Beaulieu's only claims to fame were its dry white wine, its silk shirts, one champion in footraces (Louis Garret), fourth place in the 1936 Olympics in Berlin, and, more recently, its Italian pasta factory. What earned Beaulieu this great honour? Happenstance, quite frankly, because the Prime Minister, tired of his usual routes that always pivoted around the same cities, wanted to discover what he calls "the deep country".

Beaulieu's mayor – Marc Blanc – is very worried, however. The Prime Minister's ratings haven't stopped falling since the elections. What's more, how can the protests that have a tendency to multiply during official visits be avoided? The steep cliff of Mont Saint-Pierre that leads to the village is blockaded each time the opponents on any side show their anger. On the other hand, during each of the Prime Minister's visits, the government gets into contact with the nearest prefecture and ensures that everything is done to protect him. However, a large police detachment, as we saw in Jonquière, and identity checks threaten to cause an explosion. A young member of the opposition said: "In the area, people are jealous of our freedom. If people have to show their credentials to circulate, we do not respond to the reaction of the country's people. We have the support of the entire village." Additionally, articles that appeared in La Dépêche du Centre, L'Express, Ouest Liberté and Le Nouvel Observateur suggest that activists in nearby villages intend to make the Prime Minister's day difficult. Some fanatics are even said to have undertaken a prolonged fast in the church in Saint Martinville.

The sympathetic mayor of Beaulieu does not know where to turn. He has the feeling that he finds himself in a stupid impasse. He has decided to write to the Prime Minister out of desperation to see if his village is truly a necessary stop in the planned trip. Beaulieu would rather be unknown and quiet instead of finding itself in the middle of a political battle that, through television, would be witnessed by millions of voters.

Appendix B

Model outputs for maximum pitch. P-values were calculated using the Satterthwaite approximation.

	Estimate	Std. Error	df	t-value	p-value	
(Intercept)	-1.1343	0.2724	111.4	-4.164	0.0001	***
AP vs. IP	1.1817	0.2105	99.7	5.612	<0.0001	***
AP&IP vs. Comma	-0.8607	0.1746	93.0	-4.929	<0.0001	***
Base-final penult	0.3047	0.1677	109.4	2.04	0.0464	*
Closed final syllable	-0.1440	0.3668	196.6	-2.631	0.0139	*
Closed penult	1.1533	0.1812	137.3	7.215	<0.0001	***
Heavy final vowel	-0.1859	0.3606	200.8	-0.515	0.6070	
Heavy penult vowel	0.2493	0.2749	91.5	1.817	0.0720	.
Superheavy final syllable	-0.1119	0.6721	209.6	-0.366	0.7146	

Appendix C

Model outputs for maximum amplitude. P-values were calculated using the Satterthwaite approximation.

	Estimate	Std. Error	df	t-value	p-value	
(Intercept)	1.3877	0.6574	111.4	2.111	0.0370	*
AP vs. IP	1.9705	0.5863	99.7	3.361	0.0011	**
AP&IP vs. Comma	-0.3995	0.4927	93.0	-0.811	0.4196	
Base-final penult	0.4822	0.4533	109.4	1.064	0.2900	
Closed final syllable	-3.1101	0.4576	196.6	-7.31	<0.0001	***
Closed penult	2.0582	0.7670	137.3	2.683	0.0084	**
Heavy final vowel	-0.3573	0.8689	200.8	-0.411	0.6813	
Heavy penult vowel	0.9189	0.6288	91.5	1.461	0.1483	
Superheavy final syllable	4.8909	1.4404	209.6	3.396	0.0007	***

Appendix D

Model outputs for rhyme duration. P-values were calculated using the Satterthwaite approximation.

	Estimate	Std. Error	df	t-value	p-value	
(Intercept)	-0.6097	0.0696	111.8	-8.763	<0.0001	***
AP vs. IP	-0.1068	0.0570	100.0	-1.873	0.0480	*
AP&IP vs. Comma	-0.0387	0.0479	93.3	-0.808	0.4215	
Base-final penult	-0.2454	0.0438	109.7	-3.317	0.0012	**
Closed final syllable	-0.9202	0.0911	197.2	-10.101	<0.0001	***
Closed penult	0.6992	0.0710	137.7	9.845	<0.0001	***
Heavy final vowel	-0.2662	0.0864	201.4	-2.376	0.0184	*
Heavy penult vowel	0.1881	0.0602	91.8	3.126	0.0024	**
Superheavy final syllable	0.4866	0.1412	210.2	3.447	0.0006	***

Preface to Chapter 3

Like all languages, Laurentian French is characterised by phenomena that are highly variable and, to this date, largely unexplored. Recent developments in the availability of corpora and of corpus tools enable the examination of these phenomena. Chapter 2 investigated variable prominence assignment in Laurentian French. The overarching hypothesis that weight sensitivity motivates prominence shift was tested. This hypothesis was supported by results from an analysis of corpus data. The results further supported the secondary hypotheses that vowel weight is contrastive and that codas bear weight. The implications of the results were discussed, in particular how these results suggest that Laurentian French exhibits a pitch accent that is variably attracted to a syllable bearing lexical prominence, the latter of which is computed from weight. These two types of prominence are additionally associated with different acoustic profiles: speakers modulate pitch to mark the location of the pitch accent, while speakers modulate duration and amplitude to mark the lexically prominent syllable.

Chapter 2 left open the potential interaction between processes that may help to cue prominence, and the capacity for learners to learn the structure underlying highly variable phenomena. Chapter 3 addresses these questions through an examination of the tense/lax quality of high vowels in Laurentian French. In final syllables, the tense/lax quality is fully predictable. In non-final vowels, however, the literature has proposed multiple – typically optional – phonological processes to account for the considerable variability in which high-vowel allophones surface. Furthermore, tense and lax allophones of high vowels are acoustically too similar to be distinguished using only one or two acoustic cues, which has prevented researchers from conducting large-scale analyses of production. Chapter 3 repurposes a forced aligner to perform automated classification of high vowels as tense or lax, thereby enabling empirical testing of phonological processes in production. Chapter 3 tests for phonological processes proposed in the literature to identify the grammar that a learner can generate from community-level data.

Chapter 3

Tense/Lax quality in Laurentian French

3.1 Introduction

Laurentian French (also commonly called Canadian French, Quebec French or Québécois) is characterized by a complex combination of processes affecting the tense/lax quality of high vowels. Laxing in word-final syllables is completely predictable, but laxing in non-final syllables involves both optionality and opaque interactions. While these processes have been described in detail (e.g. Dumas, 1983, 1987; Poliquin, 2006) and while considerable theoretical work has attempted to capture the patterns (e.g. Fast, 2008; Poliquin, 2006; Bosworth, 2011), no quantitative data are yet available to confirm the speaker production patterns assumed in the literature. Instead, recent descriptions and analyses generally rely on the perceptual acceptability judgments that Poliquin (2006) collected from 12 native speakers. It is essential to perform a large-scale analysis of production, however, because the claimed combination of complex interactions and considerable variability (within and across speakers) suggests that speaker acceptability judgments may not be a reliable proxy for their production patterns. An additional challenge faced by previous analyses based on Poliquin's work is that he typically abstracted away from speakers accepting multiple forms, some of which he had deemed to be ill-formed.

In the current study, we address the lack of production data directly by classifying the tense/lax quality of nearly 25,000 tokens of high vowels in non-final syllables, drawing on approximately 120 hours of speech. Previous work has demonstrated that the tense/lax quality of high vowels in Laurentian French cannot effectively be classified using only one or two acoustic dimensions (notably the first two formants; Arnaud et al., 2011; Sigouin, 2013). Furthermore, speakers vary in their articulatory implementation of the distinction (Dalton, 2011; Mielke and Dalton, 2013) and therefore classification procedures likely require a method to adapt classification to individuals. As a result, neither automatic nor manual coding of high vowels is

feasible to produce a large-scale corpus of laxing in Laurentian French. To resolve this problem, the current study employs automated classification using forced alignment, where the aligner's language models are trained on final-syllable high vowels (where tense/lax realisations are fully predictable) to enable the classification of non-final high vowels and thereby allow us to circumvent the classification difficulties described in the literature.

Consequently, we are able to provide a quantitative analysis of the community-level predictors that learners could draw upon to develop their grammar. This is a crucial step in the analysis of the tense/lax quality in Laurentian French for two main reasons. First, Poliquin (2006) proposed that learners lack sufficient input to generate a grammar of the tense/lax quality, and therefore assumed that the patterns suggested by native speakers' acceptance rates in perception result from the effects of universal constraints determining possible surface forms. With a large-scale study, we can directly test whether the input learners receive is likely to be sufficient to generate a grammar of the tense/lax quality. Second, a considerable number of optional and interacting category-changing phonological processes have been proposed to coexist in Laurentian French in an effort to capture the tense/lax quality of high vowels (e.g. laxing harmony, laxing disharmony, closed-syllable laxing, retensing after resyllabification). With a large-scale study, we can test (a) whether speaker productions support the phonological processes that have been proposed in the literature, and (b) whether laxing patterns are better captured through category-changing phonological processes or through gradient phonetic tendencies (coarticulation, lenition).

This study refutes Poliquin's claim that learners lack sufficient input to generate a grammar of the tense/lax quality in high vowels. Through employing a large corpus of production data, it additionally uncovers the grammar that a speaker is expected to infer from community-wide patterns. In so doing, we argue that category-changing phonological processes predominantly determine the tense/lax quality of high vowels, but that phonetic pressures do contribute to variability in high-vowel realisations.

3.2 The tense/lax quality of high vowels in Laurentian French

In this section, we provide an overview of the phonological processes that have been proposed to affect the tense/lax quality (henceforth *tenseness*) of high vowels in Laurentian French.[25] To contextualise high-vowel tenseness within the vowel inventory of Laurentian French, Table 5 shows both the oral and nasal monophthongs present in the dialect. We see that Laurentian French contrasts four vowel heights: high, mid-high, mid-low, and low. We assume that the distinction between mid-high and mid-low vowels is one of height rather than tenseness, following convention for French (for review and discussion, see Lamontagne, 2014), meaning that the tense/lax distinction is only relevant for high vowels.[26] While schwa is included in Figure 1, schwa has a low realisation rate and a highly restricted distribution in Laurentian French (Côté, 2008, 2012). As such, forms in which schwas introduce confounds that will be excluded from the current study (see section 3.3).[27]

Table 5: The vowel inventory of Laurentian French, adapted from Côté (2012).

Oral Monophthongs				Nasal Monophthongs		
i	y		u			
e	ø		o			
ɛ, ɛː	œ	(ə)	ɔ	ɛ̃	œ̃	ɔ̃
a			ɑ			ɑ̃

[25] Authors typically assume that high vowels are underlyingly tense in Laurentian French based on cross-dialectal comparisons and based on final-syllable laxing being predictable (e.g. Walker, 1984; but cf. Hannahs, 1989, who proposes lax high vowels as underlying). The underlying tense/lax specification of the high vowel series is not crucial to the current analysis, and we will therefore follow convention in assuming that high vowels are underlyingly tense.

[26] See section 6.2.1 for further discussion, where we revisit the assumption that the tense/lax quality is only relevant for high vowels.

[27] Our transcriptions reflect pronunciations in Laurentian French, unless stated otherwise. We omit schwas from both underlying and surface forms where they are not typically present in surface forms (for example, in word-final position; e.g. Côté, 2012). Furthermore, the symbol used for the underlying rhotic is /r/, because there is considerable variation in the realisation of this sound across dialects. It is transcribed as [ʁ] in surface forms because this is the acoustic realisation that is most common in Laurentian French (Côté and Saint-Amant Lamy, 2012). Additionally, certain non-high vowels diphthongise in Laurentian French, but our transcriptions do not reflect diphthongisation because the quality of diphthongs is subject to considerable variation that is not relevant for the current analysis (e.g. Dumas, 1974; Walker, 1984). Finally, we use [ʦ] and [ʤ] to transcribe the affrication of coronal stops, which consistently occurs before high-front vowels and glides in Laurentian French (e.g. Walker, 1984; Côté, 2012).

The following subsections focus only on high vowels. Sections 3.2.1 and 3.2.2 describe the tenseness of high vowels in final syllables, while sections 3.2.3 through 3.2.5 focus on tenseness in non-final syllables. We will largely draw examples from Poliquin's (2006) work, as he offers the most comprehensive discussion of these processes to date. The examples drawn from Poliquin will we supplemented with forms we introduce, marked with a final asterisk.

3.2.1 Lengthening in closed final syllables

Across varieties of French, vowels are lengthened when they immediately precede one of the following *lengthening consonants* in word-final position: /v z ʒ r vr/ (e.g. Grammont, 1914; Walker, 1984).[28] Illustrative examples focusing on high vowels are provided in (1). These examples provide transcriptions for normative ("standard") French in anticipation of interactions with the laxing process in Laurentian French described in section 3.2.2. Crucially, we observe that not all final codas create lengthening contexts.

(1) Lengthening (Normative French)*

a. Final Open syllable	b. Lengthening Consonant	c. Other Final Coda
/ri/ [ʁi] *rit* 'laughs'	/rir/ [ʁiːʁ] *rire* 'laugh'	/rit/ [ʁit] *rite* 'rite'
/ry/ [ʁy] *rue* 'street'	/ryz/ [ʁyːz] *ruse* 'ruse'	/rys/ [ʁys] *russe* 'Russian'
/ru/ [ʁu] *roue* 'wheel'	/ruʒ/ [ʁuːʒ] *rouge* 'red'	/rul/ [ʁul] *roule* 'rolls'

3.2.2 Laxing in closed final syllables

The tenseness of non-lengthened high vowels is fully predictable in final syllables in Laurentian French (e.g. Légaré, 1978; Dumas, 1983; Walker, 1984; Poliquin, 2006; Côté, 2012; Sigouin, 2013; Sigouin and Arnaud, 2015).[29] In final syllables closed by any coda other than a Lengthening Consonant, high vowels are obligatorally lax; in final open syllables, high vowels are obligatorily tense. This complementary distribution can be observed in morphologically related words, as illustrated in (2):

[28] We follow convention in referring to these sequences as "lengthening consonants" even though /vr/ is a consonant cluster. The lengthening consonants include voiced fricatives in addition to the rhotic.

[29] In this study, we focus on the native lexicon. In recent borrowings, laxing is not adapted and instead the source language's tense/lax quality is preserved for high vowels, like in [sut] for En. 'suit' (Légaré, 1978; Côté, 2012). Borrowings like these were manually identified and excluded from the current study.

(2) Laxing in Morphologically Related Words*

a. Tense High Vowel	b. Lax High Vowel
/p(ə)ti/ [ptsi] *petit* 'small.MASC'	/p(ə)tit/ [ptsɪt] *petite* 'small.FEM'
/mine/ [mine] *miner* 'mine.INF'	/min/ [mɪn] *mine* 'mine (n.)'
/deby/ [deby] *début* 'start (n.)'	/debyt/ [debʏt] *débute* 'start.3SG.IND.PRS'
/aʒu/ [aʒu] *ajout* 'addition'	/aʒut/ [aʒʊt] *ajoute* 'add.3SG.IND.PRS'

Vowels lengthened by a following Lengthening Consonant have been at the centre of some debate regarding their tenseness: they have been described both as tense (Légaré, 1978; Poliquin, 2006) and as lax (Cuerrier and Reiss, 2018). Côté (2010, 2012), however, argues that the pattern is slightly more complex: /r/ triggers laxing, while the other Lengthening Consonants do not. Based on a preliminary verification of the corpus under examination in this study, we follow Côté in treating vowels in final syllables closed by a Lengthening Consonant as lax only if the coda is /r/. The laxing process and its interaction with lengthening is illustrated in (3): vowels are long before /r/ like before other Lengthening Consonants (3b), but before /r/ the vowel is lax (like the vowels in (3c)), rather than tense (like in (3a)). Drawing on these data, we infer two relevant processes: *Final Closed-Syllable Laxing* and *Final Lengthening*.

(3) Closed-Syllable Laxing and Lengthening (Laurentian French)*

a. Final Open syllable	b. Lengthening Consonant	c. Other Final Coda
/ri/ [ʁi] *rit* 'laughs'	/rir/ [ʁɪːʁ] *rire* 'laugh'	/rit/ [ʁɪt] *rite* 'rite'
/ry/ [ʁy] *rue* 'street'	/ryz/ [ʁyːz] *ruse* 'ruse'	/rys/ [ʁʏs] *russe* 'Russian'
/ru/ [ʁu] *roue* 'wheel'	/ruʒ/ [ʁuːʒ] *rouge* 'red'	/rul/ [ʁʊl] *roule* 'rolls'

3.2.3 Laxing in closed non-final syllables

While high vowels obligatorily lax in final closed syllables, closed-syllable laxing is optional in non-final syllables in Laurentian French. This optional *Closed-Syllable Laxing* is illustrated in (4), where we can see that high vowels in closed syllables can surface as lax, but may also remain tense. The final syllable's segmental content is not a factor in this laxing process, which is a crucial difference from the processes we will discuss in sections 3.2.4 and 3.2.5.

(4) Non-final Closed-Syllable Laxing (Poliquin, 2006: 26)

/mistɛr/	[mɪs.tɛːr]~[mis.tɛːr]	*mystère* ‘mystery’
/bylgɑr/	[bʏl.gɑːr]~[byl.gɑːr]	*bulgare* ‘Bulgarian’
/muʃte/	[mʊʃ.te]~[muʃ.te]	*moucheté* ‘spotted’
/erypsjɔ̃/	[e.rʏp.sjɔ̃]~[e.ryp.sjɔ̃]	*éruption* ‘eruption’*
/pyrgatwɑr/	[pʏr.ga.twɑːr]~[pyr.ga.twɑːr]	*purgatoire* ‘purgatoire’*

3.2.4 Laxing Disharmony

Many authors have observed that in Laurentian French, a high vowel in the first syllable of a disyllabic word can optionally lax if it is in an open syllable and the final syllable has the same phoneme also in an open syllable (Dumas, 1983, 1987; Déchaîne, 1990; Baronian, 2001; Poliquin, 2006). In other words, a disyllabic word where both vowels are high and match in frontness and rounding allows laxing in its first syllable, but not in its second. This is illustrated in (5). While Poliquin (2006) found that most speakers accept pronunciations with laxing in the first syllable of these words, he further showed that most do not accept pronunciations with laxing in similar words with non-identical phonemes (e.g. *[ʒʏli] is rarely accepted as a pronunciation of /ʒyli/ *Julie* ‘Julie’).

(5) *Disharmony in Disyllables* (Poliquin, 2006: 160)

/midi/	[mɪʤi]~[miʤi]	*midi* ‘noon’
/isi/	[ɪsi]~[isi]	*ici* ‘here’
/zulu/	[zʊlu]~[zulu]	*zoulou* ‘Zulu’

3.2.5 Laxing Harmony

Laxing Harmony yields lax high vowels in non-final syllables in Laurentian French. This process, like disharmony, is optional. We follow Poliquin’s (2006) description of possible harmony patterns, which he based on Dumas’s (1983, 1987) earlier description as well as on acceptability judgments from his own study and on typology. Dumas suggested that different Laxing Harmony patterns are attested and may be sociolinguistically coded, although Poliquin showed that the same patterns of variation are found in speakers from a single region and age group. It is important to note that, while the patterns discussed below are described as the only options attested, native speakers in Poliquin’s study did sometimes accept other forms, including laxing without a harmonic trigger, which were usually interpreted as errors (Poliquin, 2006: 57,

77, 78), though in one additional case a judgment was included only after the stimulus was repeated by Poliquin upon receiving a response that he had not predicted (p. 152). We will not focus on these discrepancies in his proposed patterns, though we believe they should be noted. In addition, while he categorises individuals as having a specific harmony pattern, speakers frequently accepted forms that reflect a different harmony pattern.

3.2.5.1 Harmony patterns

Laxing Harmony in Laurentian French is parasitic on height: a lax vowel in the final syllable can trigger laxing of a non-final vowel only if both vowels are high. For instance, while /mitik/ *mythique* 'mythical' allows both a lax ([mɪtsɪk]) or a tense ([mitsɪk]) initial vowel, /mitɛn/ *mitaine* 'mitten' must be realised with a tense initial vowel ([mitɛn], *[mɪtɛn]). Examples of Laxing Harmony are provided in (6):

(6) Laxing Harmony in Disyllables (Poliquin, 2006: 54)

/mitik/	[mɪtsɪk]~[mitsɪk]	*mythique* 'mythical'
/minyt/	[mɪnʏt] ~[minʏt]	*minutes* 'minutes'
/stypid/	[stsʏpɪd] ~[stsypɪd]	*stupide* 'stupid'
/kutym/	[kʊtsʏm] ~[kutsʏm]	*coutûmes* 'customs'

In principle, disyllabic forms are ambiguous between two potential target types for harmony. Laxing Harmony could target either the word's initial syllable or the syllable immediately preceding a laxing trigger. Furthermore, the latter pattern could be iterative or not. In fact, all of these options are attested in longer words. Accordingly, following Poliquin (2006), we will refer to the pattern where harmony targets the initial syllable only as Non-Local Harmony and the pattern where harmony targets the immediately preceding vowel as Local Harmony, whether it is iterative or not. Poliquin (2006) and Dumas (1987) suggested that individual speakers differ in their dominant harmony patterns, though Laxing Harmony is an optional process for all speakers who have this process.

As mentioned, in words with at least three syllables, different harmony processes result in distinct surface forms, as we illustrate in Table 1.[30] When no harmony occurs (row a), a word like /ilymin/ *illumine* 'illuminates' or /similityd/ *similitude* 'similarity' will have a lax high vowel in the final syllable (obligatory Closed-Syllable Laxing in final syllables) but tense high vowels elsewhere because no harmony process motivates laxing in non-final syllables and all non-final syllables are open. When local harmony applies, the result is [ilʏmɪn] and [similɪtsʏd] if harmony is non-iterative (row b); if it is iterative, the target of each application of harmony (e.g. the penult) becomes the trigger for harmony in the preceding syllable (e.g. the antepenult), yielding [ɪlʏmɪn] and [sɪmɪlɪtsʏd] (row c). Non-Local Harmony (row d) differs from local harmony in that the initial syllable is targeted regardless of intervening content, which Poliquin assumes is because non-local harmony targets the syllable bearing secondary stress.[31] This generates [ɪlymɪn] and [sɪmilɪtsʏd]. The remaining forms in Table 6 are expected outputs when at least two harmony patterns apply, but it is not certain from Poliquin to what extent these forms occur. If both Local Non-Iterative Harmony and Non-Local Harmony apply (row e), the resulting surface form is identical to the one resulting from Local Iterative Harmony unless the word has at least four consecutive syllables with potential targets (e.g. [sɪmilɪtsʏd] with a tense vowel in the second syllable) or if the word has medial non-high vowels (see section 3.2.5.2). When Local Iterative Harmony applies, it is not possible to determine whether Non-Local Harmony has applied as well from the surface form unless there are intervening non-high vowels, which we turn to next.

[30] Following Poliquin, this table does not include the possibility of a form having *Local Iterative Harmony* that is non-exhaustive due to this pattern being unattested typologically (e.g. the process applies twice, but does not target subsequent vowels). We follow Poliquin's assumption that *Non-Local Harmony* targets the initial syllable and therefore that the second syllable in *similitude* cannot be the only non-final syllable that undergoes laxing. These forms correspond to patterns given non-null likelihood in Fast's (2008) computation of possible surface forms in her Optimality Theoretic analysis, with the exception of [sɪmɪlitsʏd], which can occur only if local harmony applies left-to-right. We will return to the possibility of forms like this in section 6.2.2 of the Discussion.

[31] Poliquin (2006) notes that some speakers accept forms where *Non-Local Harmony* targets the leftmost high vowel, whether that vowel is initial or not, such as in /definitif/ *définitif* 'final' pronounced [definɪtsɪf], but we will not focus on this possibility in the current analysis (though our models indirectly offer some ability to test this). In these cases, it is not clear why a high vowel in a non-initial syllable would be targeted by laxing harmony following Poliquin's explanation that secondary prominence is word-initial and that Non-Local Harmony targets secondary prominence. Hannahs (1989) suggests that all high-vowel laxing in Laurentian French is sensitive to prominence, but we leave testing the particulars of his proposal for future work because his analysis differs considerably from more recent proposals.

Table 6: Possible realisations for /mitik/ mythique ‘mythical’, /ilymin/ illumine ‘illuminates’ and /similityd/ similitude ‘similarity’ based on the harmony processes that apply. For ease of reading, we have included a reference string for whether the syllable contains a tense (“T”) or lax (“L”) high vowel.

Harmony Processes	*Mythique*	*Illumine*	*Similitude*
a. No harmony	[mitsɪk], TL	[ilymɪn], TTL	[similitsʏd], TTTL
b. *Local Non-Iterative Harmony*	[mɪtsɪk], LL	[ilʏmɪn], TLL	[similɪtsʏd], TTLL
c. *Local Iterative Harmony*	[mɪtsɪk], LL	[ɪlʏmɪn], LLL	[sɪmɪlɪtsʏd], LLLL
d. *Non-Local Harmony*	[mɪtsɪk], LL	[ɪlymɪn], LTL	[sɪmilitsʏd], LTTL
e. *Local Non-Iterative Harmony* and *Non-Local Harmony*	[mɪtsɪk], LL	[ɪlʏmɪn]. LLL	[sɪmilɪtsʏd], LTLL
f. *Local Iterative Harmony* and *Non-Local Harmony*	[mɪtsɪk], LL	[ɪlʏmɪn], LLL	[sɪmɪlɪtsʏd], LLLL

3.2.5.2 Neutral vowels

As Laxing Harmony in Laurentian French is parasitic on height, non-high vowels do not participate in the process (Poliquin, 2006). When such vowels intervene between target and trigger, they thus have the status of neutral vowels. Poliquin finds that speakers differ with respect to their treatment of neutral vowels, of which he included only /e/ in his stimuli. Some speakers treat neutral vowels as transparent, in which case the vowels are invisible to harmony and therefore harmony occurs as though they were not present. Other speakers treat neutral vowels as opaque, in which case the vowels are visible to harmony and block the process from applying past them. These possibilities are illustrated in (7), where we show that opaque and transparent vowels only produce distinct surface forms when harmony is local. Local Iterative Harmony only permits laxing in the initial syllable of a word like /inedit/ *inédite* ‘unpublished (fem.)’ if neutral vowels (here /e/) are transparent (as in (7a)); otherwise the neutral vowel blocks harmony (as in (7b)) and the initial vowel remains tense. We see in (7c) that the initial vowel can undergo Laxing Harmony when there is a neutral intervening vowel when Non-Local Harmony applies, which yields surface forms for trisyllabic words that are identical to those that occur when a word undergoing Local Iterative Harmony contains a transparent medial vowel (7a).

(7) Harmony Patterns with Neutral Vowels (adapted from Poliquin, 2006: 83)

a. Local Iterative Harmony with Transparent Intervening Vowels

/inedit/	[ɪnedzɪt]~[inedzɪt]	*inédite* 'unpublished (fem.)'
/yterys/	[ʏterʏs]~[yterʏs]	*utérus* 'uterus'
/lykratif/	[lʏkratsɪf]~[lykratsɪf]	*lucratif* 'lucrative'*

b. Local Iterative Harmony with Opaque Intervening Vowels

/inedit/	[inedzɪt]	*inédite* 'unpublished (fem.)'
/yterys/	[yterʏs]	*utérus* 'uterus'
/lykratif/	[lykratsɪf]	*lucratif* 'lucrative'*

c. Non-Local Harmony with Any Neutral Intervening Vowels

/inedit/	[ɪnedzɪt]~[inedzɪt]	*inédite* 'unpublished (fem.)'
/yterys/	[ʏterʏs]~[yterʏs]	*utérus* 'uterus'
/lykratif/	[lʏkratsɪf]~[lykratsɪf]	*lucratif* 'lucrative'*

However, the medial neutral vowel is not treated the same way. With Non-Local Harmony, the neutral vowel is not computed as part of the application of Laxing Harmony, and high vowels in medial syllables will not lax. With Local Iterative Harmony, on the other hand, the neutral vowel does not block harmony, but a high vowel in its place will undergo laxing. Table 2 summarises the treatment of neutral vowels comparing the possible realisations for *inédite* 'unpublished (fem.)' with those for *illumine* 'illuminates', highlighting the differences between the harmony processes.

Table 7: Possible realisations for /inedit/ inédite 'unpublished (fem.)' and /ilymin/ illumine 'illuminates' based on the harmony processes that apply. As before, we include an abstracted transcription alongside the example, with "N" representing a neutral vowel.

Processes	*Inédite* /inedit/	*Illumine* /ilymin/
No harmony	[inedzɪt], TNL	[ilymɪn], TTL
Local Non-Iterative Harmony	[inedzɪt], TNL	[ilʏmɪn], TLL
Local Iterative Harmony	[inedzɪt] or [ɪnedzɪt], TNL or LNL	[ɪlʏmɪn], LLL
Non-Local Harmony	[ɪnedzɪt], LNL	[ɪlymɪn], LTL
Local Non-Iterative Harmony and *Non-Local Harmony*	[ɪnedzɪt], LNL	[ɪlʏmɪn], LLL
Local Iterative Harmony and *Non-Local Harmony*	[ɪnedzɪt], LNL	[ɪlʏmɪn], LLL

3.2.5.3 Geminates and Non-Local Harmony

Cuerrier and Reiss (2018) respond to Poliquin's (2006) proposal that word-initial syllables are targeted by Non-Local Harmony by offering an alternative explanation for initial-syllable laxing.

They observe that many such words contain a high-vowelled prefix. In a form like /ilisit/ *illicite* 'illicit, they argue that the /l/ is part of a (derived) geminate, and therefore that laxing in this word is due to what we will refer to as *Pre-Geminate Laxing*. They support this proposal in two ways based on speaker judgments they collect. First, producing initial laxing in *illicite*, as in [ɪl(:)isɪt], does not predict initial-syllable laxing in disyllabic words like /lisit/ *licit* 'licit', suggesting that Laxing Harmony does not explain initial-syllable laxing in *illicite*. Second, initial-syllable laxing is possible in /il(:)egal/ *illégal* 'illegal' despite the apparent lack of a trigger for laxing (lax high vowel, closed syllable). While this proposal cannot explain judgments in all words that Poliquin used as stimuli, it is a possibility that has not yet received systematic empirical confirmation.

3.2.5.4 Lengthening and Laxing Harmony

In section 3.2.1, we showed that vowels are obligatorily lengthened when in final syllables closed by a Lengthening Consonant (/v z ʒ r vr/), and in section 3.2.2, we showed that high vowels in closed syllables obligatorily undergo laxing in Laurentian French. Before /r/, high vowels are both long and lax (e.g. /vir/ [vɪːʁ] *vire* 'turns'), while high vowels lengthened by other Lengthening Consonants are both long and tense (e.g. /vivr/ [viːvʁ] *vivre* 'to live'). Surprisingly, as illustrated in (8), Laxing Harmony can occur when the surface realisation of the final-syllable high vowel is tense due to Pre-Fricative Tensing (Dumas, 1983, 1987; Poliquin, 2006), which is consistent with the final-syllable high vowel having undergone Closed-Syllable Laxing at an earlier stage of the derivation. Poliquin demonstrates that laxing in the initial syllable of /misiv/ *missive* 'letter' could not be due to Laxing Disharmony (section 3.2.4) because *Laxing Disharmony* requires identical vowels, while the examples in (8) reveal that laxing triggered by lengthened high vowels does not (e.g. /difyz/ *diffuse* 'diffuses' realised [dʑɪfyːz]).

We illustrate process interactions involving lengthening with the word /misiv/ *missive* 'letter'. First, obligatory Closed-Syllable Laxing results in a lax high vowel in the final syllable. Following this, optional Laxing Harmony can cause the high vowel in the initial syllable to lax, with the lax high vowel in the final syllable as the trigger. Finally, obligatory Pre-Fricative Tensing yields a tense high vowel in the final syllable because the syllable is closed by a voiced fricative (/v z ʒ/). Lengthening could occur at any step of this derivation so long as Lengthening

either occurs after Closed-Syllable Laxing or the Closed-Syllable Laxing rule does not require the target vowel to be short. Consequently, if Laxing Harmony applies, the output for *missive* is the opaque surface form [mɪsi:v], in which the final syllable contains a tense high vowel, even though Laxing Harmony resulted in a lax high vowel in the initial syllable. In a rule-based account, Pre-Fricative Tensing could therefore be described as *retensing*, as evidenced by its interaction with Laxing Harmony.

(8) Harmony and Pre-Fricative Tensing (Poliquin, 2006: 221)

/misiv/	[mɪsi:v]~[misi:v]	*missive* 'letter'
/difyz/	[ʤɪfy:z]~[ʤify:z]	*diffuse* 'diffuses'
/litiʒ/	[lɪtsi:ʒ]~[litsi:ʒ]	*litige* 'litigates'

3.2.5.5 Obligatory Retensing from Resyllabification

Like Pre-Fricative Tensing, resyllabification that results from affixation can result in opaque surface forms. High vowels that are followed by a consonant that has been resyllabified from coda to the onset of the next syllable are obligatorily tense in Laurentian French (Poliquin, 2006). For example, the /i/ in /myzik/ *musique* 'music' always surfaces as tense in /myzik+al/ *musical* 'musical'. We can see from the forms in (9) that Laxing Harmony can apply to the initial syllable of *musical*, producing [mʏzikal]~[myzikal], even though the following high vowel trigger surfaces as tense. This is because the trigger is at the right edge of the base and, thus, the structural description for laxing is met in this environment (/myzik/ [mʏzɪk]~[myzɪk] *musique*). From a rule-based perspective, both Closed-Syllable Laxing and Laxing Harmony applied prior to the attachment of the suffix, the latter of which triggered Retensing from Resyllabification.

(9) Harmony and *Retensing from Resyllabification* (Poliquin, 2006: 146)

Form without resyllabification	Form with resyllabification
/myzik/ [mʏ.zɪk]~[my.zɪk] *musique* 'music'	/myzik+al/ [mʏ.zi.kal]~[my.zi.kal] *musical* 'musical'
/sivil/ [sɪ.vɪl]~[si.vɪl] *civil* 'civil'	/sivil+ite/ [sɪ.vi.li.te]~[si.vi.li.te] *civilité* 'civilness'
/yrin/ [ʏ.ʁɪn]~[y.ʁɪn] *urine* 'urine'	/yrin+wɑr/ [ʏ.ʁi.nwɑʁ]~[y.ʁi.nwɑʁ] *urinoire* 'urinal'

Unlike word-internal resyllabification across morphemes, resyllabification across words (i.e. *enchaînement* 'linking') crucially does not trigger *Retensing from Resyllabification*. On one hand, the /i/ in /myzik+al/ *musical* obligatorily surfaces as tense due to *Retensing from Resyllabification* ([mʏzikal]~[myzikal]). On the other hand, the /i/ in /myzik#a/ *musique à* 'music of' obligatorily surfaces as lax due to *Closed-Syllable Laxing* ([mʏzɪka]~[myzɪka]), even if the [k] is realised as the onset of the first syllable of the following word. In both cases, the surface form resulting from these obligatory processes is fully predictable.

3.2.6 High-Vowel Tenseness and Deletion

The question of predicting tenseness in high vowels is further confounded by the possibility of deleting high vowels.[32] High-Vowel Deletion may interfere with the learner's capacity to infer a grammar from other speakers' realisations because the trigger for harmony may have been deleted or because deletion triggered the creation of a new coda (which can optionally trigger laxing). The challenge posed by High-Vowel Deletion is noted by Poliquin (2006: 24) but not tested and the effect it may have on speakers' acceptability judgments is not discussed. In Table 3, we illustrate forms that have undergone the laxing processes discussed earlier, with the possibility of laxing in closed syllables derived from High-Vowel Deletion in the surface form [dzɪf.sɪl] as proposed by Poliquin.[33] The challenge is that, when High-Vowel Deletion applies, a listener cannot determine from the surface form whether Laxing Harmony applied or Closed-Syllable Laxing applied.

[32] Furthermore, there is High-Vowel Devoicing, which is sometimes equated with High-Vowel Deletion (e.g. see Walker, 1984) despite distinct phonological conditioning (Garcia, Goad and Guzzo, 2017). We will not distinguish devoiced tokens from voiced ones, as this difference is not crucial to the analysis and the aligner does not consider voicing directly. See section 4.3 for additional information.

[33] For brevity, we do not illustrate the possibility of deleting from initial syllables, which Garcia, Goad and Guzzo (2017) show is disfavoured in Laurentian French (but see further section 4.3).

Table 8: Possible realisations for /difisil/ difficile 'difficult' based on the harmony processes that apply. As before, we include an abstracted transcription alongside the example, with "0" representing a deleted vowel.

Processes	/difisil/ *Difficile*
No harmony	[dzi.fi.sɪl], [dzif.sɪl] or [dzɪf.sɪl]; TTL/T0L/L0L
Local Non-Iterative Harmony	[dzi.fɪ.sɪl], [dzif.sɪl][34] or [dzɪf.sɪl]; TLL/T0L/L0L
Local Iterative Harmony	[dzɪ.fɪ.sɪl] or [dzɪf.sɪl]; LLL/L0L
Non-Local Harmony	[dzɪ.fi.sɪl] or [dzɪf.sɪl]; LTL/L0L
Local Non-Iterative Harmony and *Non-Local Harmony*	[dzɪ.fɪ.sɪl] or [dzɪf.sɪl]; LLL/L0L
Local Iterative Harmony and *Non-Local Harmony*	[dzɪ.fɪ.sɪl] or [dzɪf.sɪl]; LLL/L0L

3.2.7 Summary

As we have shown, non-final syllables pose a challenge for a description of the distribution of tense and lax high vowels in Laurentian French. While the tenseness of high vowels is fully predictable in final syllables, many interacting – and often optional – processes render tenseness difficult to predict in non-final syllables: optional Closed-Syllable Laxing, optional Disharmony, optional Local Harmony (iterative or not), optional Non-Local Harmony, obligatory Retensing from Resyllabification, and obligatory Pre-Fricative Retensing. The issue is further complicated by the differing treatments of neutral vowels (opaque or transparent) and by the French lexicon (morphological structure), which includes a large number of derived words and therefore multiple possible points in the derivation at which processes could have applied (assuming cyclicity, as proposed by Poliquin, 2006 and Fast, 2008). Additionally, Poliquin suggests that learners do not have sufficient input to reliably determine the type of Laxing Harmony that their grammar should include given that the type of harmony that is operative is often indeterminate in words with fewer than three syllables. Poliquin consequently proposes that speakers' grammars vary because speakers must infer a Laxing Harmony pattern without sufficient input from the community to reliably converge on a single Laxing Harmony process (trigger, target, and interativity if Local Harmony).

[34] While the other forms listed do not rely on the relative ordering of harmony and deletion, this form only arises if harmony precedes deletion.

3.3 Research questions

With the complex interplay of processes in mind, this study addresses two main research questions. The first question is predicated on laxing being the result of category-changing phonological processes:

> Which phonological processes determining high vowel tenseness can learners be expected to generate based on naturalistic input from the speech community?

In other words, based on input learners are exposed to in the speech community, we seek to determine the community-level grammar of laxing, leaving for future work a comparison of this community-level grammar with the systems that individual speakers build from that input. With this objective in mind, we probe the patterns induced by mixed-effects logistic regression, thereby treating the statistical model like a learner of Laurentian French. Following the literature, we expect to find that the following processes that apply in non-final syllables will be present in the community-level grammar: (a) Laxing Disharmony, (b) Laxing Harmony, (c) Closed-Syllable Laxing (including Pre-Geminate Laxing), and (d) Tensing from Resyllabification. With respect to Laxing Harmony, we aim to determine which types of Laxing Harmony (Local Iterative, Local Non-Interative, or Non-Local) learners would generalise from the community-level data. We expect that high vowels will be more likely to undergo laxing when there is a lax high vowel in a following syllable (Local Non-Iterative Harmony in penults; Local Iterative Harmony elsewhere) and in the final syllable (Non-Local Harmony in word-initial syllables; Local Iterative Harmony with transparent neutral vowels). Contrary to Poliquin's (2006) claim that speakers do not receive enough input from words of three or more syllables to determine the triggers and targets of Laxing Harmony, we expect that (a) the learner is capable of inferring Laxing Harmony processes directly from their input, and (b) the learner need not rely on innate principles (e.g. locality) to infer different processes from the same phonological input.

The second question pertains to the phonological status of high-vowel tenseness in Laurentian French:

> Is laxing truly governed by the operation of phonological processes in non-final syllables or does it instead result from phonetic pressures? If laxing does result from phonetic pressures, is it phonetic reduction, coarticulation, or both?

We have demonstrated that the many processes proposed to account for the laxing of high vowels in non-final syllables allow for immense intraspeaker and interspeaker variation for a wide array of words. It may thus be that a fundamental assumption in the literature has been incorrect; not all cases of high-vowel laxing result from category-changing phonological processes. For instance, laxing may result from phonetic reduction, namely articulatory undershoot that results in a high vowel that is less peripheral than its target (e.g. Lindblom, 1963; Flemming, 2005). Ultrasound studies (Dalton, 2011; Mielke and Dalton, 2013) show that speakers of Laurentian French use tongue body position to articulate a laxed high vowel (only some speakers use tongue root position in addition to tongue body position) and therefore undershoot would often result in a high vowel that would be perceived as lax; this may not have been the case if laxing were best predicted by tongue root position. Alternatively, Laxing Harmony may involve coarticulation (i.e. the process may not have been phonologised), in which case laxing results from phonetic pressures in addition to or instead of category-changing phonological processes. We expect that the community-level grammar will include phonological processes, but that these processes will be supplemented by both (a) gradient phonetic reduction, and (b) gradient vowel-on-vowel coarticulation.

Both research questions we have forwarded require two assumptions to be satisfied; otherwise, they cannot be answered. First, Poliquin (2006) must have underestimated the amount of data that learners receive as input. Otherwise, our sample will almost certainly be insufficient to generate a community-level grammar. Second, tense and lax high vowels must be sufficiently acoustically distinct – in both final and non-final positions – that they can be classified with high accuracy in order to perform a large-scale analysis of tenseness in Laurentian French. We describe our methods in the following section, where we show that both assumptions are satisfied.

3.4 Methods

The overall objective of this study is to determine the community-level patterns for the distribution of tense and lax high vowels in Laurentian French. Section 3.4.1 provides a description of the corpus used in the study. Sections 3.4.2 and 3.4.3 then describe the forced aligner and the procedure used to classify the surface realisations of high vowels as tense or lax,

respectively. Section 3.4.4 then provides a description of the statistical analysis used in this study, including predictors and how they were coded. Finally, section 3.4.5 summarises the methodology and predictors.

3.4.1 Corpus

The corpus used in this study is drawn from two sources. The first is political debates from the *Assemblée nationale du Québec* ('Quebec National Assembly'), whose website provides audio recordings and written transcripts (ANQ, 2011; see Milne, 2014). The debates are from May 2011, and contain approximately 60 hours of speech from 105 different speakers. While all speakers provided training data, namely, high vowels in final syllables (see section 3.4.3), 45 spoke so little that they did not provide any words with high vowels in non-final syllables. Only the remaining 60 speakers from this first corpus are thus included in our analysis of tenseness in non-final syllables. The second corpus is compiled from 8 survey locations in the *Phonologie du français contemporain* corpus (PFC; Durand, Laks and Lyche, 2002; https://www.projet-pfc.net/), providing data from the Canadian provinces of Quebec (Chelsey, Chicoutimi, La Pocatière, Pointe-Fortune, Quebec City, Trois-Rivières and Vanier) and Ontario (Hawkesbury). In total, the PFC contributed approximately 60 hours of speech from 71 speakers. Data from both sources were grouped together for the analysis, although the conversational context differs between – and often within – the source corpora, creating a corpus of about 120 hours of speech.

3.4.2 Forced Aligner

We adapt the SPLAlign forced aligner (Milne, 2014) for the present study. During forced alignment, the aligner is provided four essential components: (1) an audio file to align; (2) an orthographic transcription of the audio to be aligned; (3) a pronunciation dictionary with words and their pronunciations to convert the orthographic transcription into a transcription conveying speech sounds (phones); and (4) language models that the aligner uses in order to match phones in the audio using spectral properties of the corresponding phones. The output of forced alignment is a segmented and annotated audio file, such as a Praat Textgrid, which includes words and phones. Crucially for this study, the aligner can be provided with more than one

possible phone sequence (i.e. possible pronunciation), in which case the aligner determines which pronunciation best matches the audio it is aligning and provides boundaries and phone annotations based on that match.

The SPLAlign forced aligner boasts several properties that make it ideal for the current study. First, the aligner includes speaker adaptation; the aligner's language models adapt to the acoustic realisations of the speaker whose audio is being aligned and, as a result, it achieves higher accuracy rates for each speaker because the aligner can better accommodate individual differences. Given that we know speakers resort to different articulatory implementations of the tense/lax distinction in high vowels (Dalton, 2011; Dalton and Mielke, 2013), this enables the aligner to produce a more reliable classification of this distinction in high vowels. The aligner offers a second benefit: its use of cross-word triphone models. When aligning an audio file, the aligner takes into account phonological context by having language models that factor in the preceding and following segments (hence triphones). For example, when aligning the /t/ in /sity/ *situe* 'situates', the aligner uses a language model for /t/ that was generated specifically from tokens of /t/ preceding /y/ and following /i/. As a result, the aligner will expect and include affrication in the /t/ in *situe* anticipating [sitsy], but not in the /t/ in /site/ *cité* 'cited', which is pronounced [site]. Furthermore, the aligner ignores word boundaries when determining the segments that flank a sound as long as no pause intervenes, because alignment comparisons showed this to yield better results for French (see Milne, 2014), likely due to the application of cross-word resyllabification increasing coarticulatory effects between segments in adjacent words. For example, in /si#te/ *si t'es* 'if you are', the aligner will consider the /t/ as being between /i/ and /e/ just as it does in *cité*. This is beneficial for the current study because tense and lax high vowels are in complementary distribution in final syllables: the onset of a following word allows the aligner to include predictable coarticulatory effects in its models for tense high vowels (e.g. a tense [i] before /t/ in /si#te/ *si t'es* 'if you are') and to include tokens of lax high vowels in open syllables (e.g. a lax [ɪ] before an onset /t/ in /sit#e/ *site est* 'site is'). The use of cross-word language models therefore enables the aligner to learn predictable coarticulatory influences of immediately adjacent phones on high vowels, particularly for high tense vowels.

As we noted in section 3.3, high vowels may undergo deletion or devoicing in non-final syllables, a phenomenon that the SPLAlign forced aligner can accommodate. In the case of

deletion, the aligner generates a pronunciation with deletion (e.g. [vzaʒ] for /vizaʒ/ *visage* 'face') and without (e.g. [vizaʒ]), and determines which sequence of segments is most likely given the audio. Because the SPLAlign forced aligner's language models are tristate, the aligner attempts to match three consecutive slices for each segment, which in a vowel roughly correspond to the onset, steady state and offset. If the aligner cannot match these three slices (a minimum of thirty milliseconds, since each slice is ten milliseconds), then the aligner does not align the segment as being present. While this is designed to capture segment-inherent spectral changes as well as coarticulation, for the current study, it gives a point of reference for when the aligner will classify a segment as being deleted. Additionally, the aligner strips the pitch track from the audio file because it was found to reduce alignment accuracy for Laurentian French, likely due to devoicing in the dialect. As a result, the aligner has been trained on voiced and devoiced tokens, and selects the phone to annotate without directly factoring in the pitch track.

3.4.3 Classifying high vowels

As described in section 3.1, tenseness is difficult to classify automatically in Laurentian French high vowels due to the considerable acoustic similarity between the tense and lax variants of a given phoneme. In this study, the use of a forced aligner provides a solution to this problem: high vowels are categorised not by one or two acoustic cues (e.g. the first two formants), but instead by the broader spectral profile of a vowel (i.e. the full spectrum, other than pitch, with three temporally ordered acoustic targets to capture dynamic properties). The aligner incorporates information about the predictable coarticulatory effects of immediately adjacent segments, and additionally adapts the alignment to speaker-specific differences in spectral profiles and coarticulation.

To train the aligner, we use the final-syllable tokens of high vowels because their realisations are predictable; the high vowels in closed syllables will be lax (unless closed by a lengthening consonant other than [ʁ]), and those in open syllables will be tense. In total, there were 48,357 tokens of final-syllable high vowels, and the phone models were trained on 80% of these while reserving the remaining 20% of tokens to test the accuracy of the aligner's classification. The aligner achieved 87.56% accuracy on the test items treating the predicted realisation as correct (i.e. a lax high vowel in final closed syllables and a tense high vowel in final open syllables). The

aligner's accuracy was further verified by partitioning the data into subsamples that alternate in their function as training and test data (k-fold cross-validation) so that each token alternates between being in the training data and being in the test data to vary accuracy across the whole dataset, yielding accuracy rates of up to 91%.

With high accuracy levels achieved for classification, the aligner was run to classify high vowels in non-final syllables. When aligning words containing non-final high vowels, the aligner generated three possible pronunciations for each of those high vowels: the vowel output could be tense, lax, or deleted.[35] Table 9 shows possible pronunciations for the first two vowels in the word *illumine* 'illuminates' based on the possible alignment outputs. We deemed it important to include the possibility of deletion, which as noted in the previous section will occur whenever the aligner cannot assign at least three ten-millisecond windows to the target vowel.[36] The aligner can thus independently classify each token, without being restricted to surface forms that might be expected a priori.

Table 9: Possible aligner outputs for /ilymin/ illumine 'illuminate'.

Aligner output	Vowel 1	Vowel 2
ilymɪn	Tense	Tense
ilʏmɪn	Tense	Laxed
ilmɪn	Tense	Deleted
ɪlymɪn	Laxed	Tense
ɪlʏmɪn	Laxed	Laxed
ɪlmɪn	Laxed	Deleted
lymɪn	Deleted	Tense
lʏmɪn	Deleted	Laxed
lmɪn	Deleted	Deleted

Following alignment, there were 24,631 words with (non-deleted) high vowels in non-final syllables, with a total of 3,828 unique lexical items. We excluded 693 tokens in which a schwa was pronounced in a potential harmony trigger syllable (based on our factors, see section 3.4.4.1)

[35] Our analysis of gradient effects due to phonetic pressures will examine the gradient likelihood of a vowel being classified as tense or lax, but the vowel's classification remains categorical, as tense or lax. To simplify, if the aligner determines the signal is an 85% match for [i] and a 40% match for [ɪ], the aligner will align the signal as containing [i]. See sections 4.4.1 and 4.4.5 for details on the statistical models used to test phonetic effects.

[36] We initially plotted data and tested models that included the presence of deletion as a predictor for bisyllabic tokens, but it was excluded from the final models because it was not a significant predictor in test models (β=0.13660, p=0.5824). As such, laxing in the initial syllable of /difisil/ *difficile* 'difficile' patterned with the initial syllable of disyllabic words when the medial vowel was deleted, mirroring a word like /mistik/ *mystique* 'mystical'.

or in any syllable following a high vowel. For example a token of /min/ *mine* 'mine' would be excluded if it were pronounced [mɪnə], or /fasilmɑ̃/ *facilement* 'easily' would be excluded if it were pronounced [fasɪləmɑ̃] or [fasiləmɑ̃]. We excluded tokens with schwas because closed-syllable processes can be triggered even when the consonants preceding schwa are seemingly in a following onset (which confounds Closed-Syllable Laxing). For example, /min/ *mine* 'mine' regularly undergoes laxing whether a schwa is realised ([mɪnə]) or not ([mɪn]), even though the high vowel appears to be neither in a final syllable nor in a closed syllable in the first surface form. This suggests that syllabification does not treat schwas and other vowels in the same manner (for discussion on the behaviour of French schwa in syllabification, see Noske, 1982). As such, we limit ourselves to the remaining 23,938 tokens and we leave for future work an examination of how schwas interact with processes affecting high-vowel tenseness.

3.4.4 Statistical analysis

We perform mixed-effects logistic regression to predict the tenseness of high vowels in non-final syllables using the glmer function in lme4 (Bates et al., 2015). We include random intercepts for words and speakers because both may show idiosyncratic preferences for a syllable to be laxed or not for reasons not captured by the current analysis. The dependent variable in each model was the high vowel in a given position (e.g. the initial syllable) within one of three types of words (disyllabic, trisyllabic, or longer words). This yields 6 models: initial syllables for each type of word, penults for trisyllabic and longer words, and medial (between the initial syllable and the penult) in longer words. The benefit of using logistic regression for each syllable rather than multinomial regression with each possible word is that it facilitates the interpretation of how vowels influence each other (e.g. distinguishing the effects of the following vowel and the final vowel). The aligner is provided no built-in biases in how harmony is expected to apply, and instead solely determines the effects of our predictors.

We will provide the number of tokens, the number of unique words, and the number of speakers for each model in section 3.5, alongside the description of the results from that model. In sections 3.4.4.1 through 3.4.4.5, we discuss the factors included in the analysis and how those factors were coded. The factors were selected to test the effects of other vowels from the same word (Laxing Harmony, coarticulation), of syllable structure (Closed-Syllable Laxing, potential

geminates), of Lengthening Consonants (Pre-Fricative Tensing, Lengthening Consonants), and of morphology (Retensing from Resyllabification, opaque Laxing Harmony). As we are interested in addressing the phonological status of laxing (see question 2 in section 3.3), we include speech rate as a predictor to test whether reduction or coarticulation can explain the laxing patterns observed.[37] The factors and the motivations for including those factors are summarised in Table 10.

Table 10: Factors tested in the current study.

Factor	Section	Phenomena Tested
Potential conditioning vowels	3.4.4.1	Laxing Disharmony (§3.2.4), Laxing Harmony (§3.2.5), Gradient coarticulation (§3.3)
Syllable structure	3.4.4.2	Closed-Syllable Laxing (§3.2.3), Pre-Geminate Tensing (§3.2.5.3)
Lengthening Consonants	3.4.4.3	Pre-Fricative Tensing (§3.2.1)
Morphological structure	3.4.4.4	Laxing Harmony (§3.2.5), Retensing from Resyllabification (§3.2.5.5)
Morphological structure x Potential conditioning vowels	3.4.4.4	Laxing Disharmony (§3.2.4)
Morphological structure x Syllable structure	3.4.4.4	Closed-Syllable Laxing (§3.2.3)
Morphological structure x Lengthening Consonants	3.4.4.4	Pre-Fricative Tensing (§3.2.1)
Speech Rate	3.4.4.5	Phonetic reduction (§3.3)
Speech Rate x Potential conditioning vowels	3.4.4.5	Gradient coarticulation (§3.3)

3.4.4.1 Potential conditioning vowels

As one of our goals is to establish the extent to which vowel harmony explains laxing and the extent to which the trigger for that harmony might be evident at the community level, we code the vowel height (and, for high vowels, the vowel tenseness) of potential triggers for harmony or sources of coarticulation. We refer to these vowels as *conditioning vowels* because on one hand the vowel's quality could predict the application of harmony, disharmony or coarticulation, but on the other hand they may not be a trigger for phonological harmony or disharmony (e.g. /a/ is a

[37] As discussed in chapter 4, lexical frequency is a predictor commonly associated with phonetic reduction and phonological variability. However, Poliquin (2006) suggested that frequency is not associated with laxing rates – admittedly based on a limited sample and perceptual acceptability judgments rather than production data. As such, given the relatively large number of predictors of interest in this study and suggestive evidence that frequency is not associated with laxing, we leave the investigation of frequency effects on high-vowel laxing for future work.

conditioning vowel if it is in a syllable with the potential to influence high-vowel laxing). We therefore code up to four vowels based on their positions relative to the vowel being analysed in the statistical model: (a) the initial syllable's vowel, (b) the preceding syllable's vowel, (c) the following syllable's vowel, and (d) the final syllable's vowel. For example, when trying to predict the tenseness of medial /i/ in /ɛ̃sybɔrdinasjɔ̃/ *insubordination* 'insubordination', the model includes that the initial syllable's vowel is a mid-low vowel (/ɛ̃/), that the preceding syllable's vowel is a mid-low vowel (/ɔ/), that the following syllable's vowel is a low vowel (/a/) and that final syllable's vowel is a mid-low vowel (/ɔ̃/). If two factors consistently refer to the same vowel in a model, then only one of the two factors is included in the model. In the model analysing the penult of trisyllabic words like /ilymin/ *illumine* 'illuminates', for instance, the initial syllable is also the preceding syllable. Consequently, the model predicting high-vowel tenseness in the penult of trisyllabic words only includes one of the two factors.

The inferred harmony pattern varies based on the position of vowels that trigger laxing and therefore by factor in the model. Bearing in mind the possibility of bi-directional or left-to-right Laxing Harmony, which is raised by Poliquin (2006), we consider the quality of conditioning vowels on either side of the target vowel rather than only vowels later in the word (i.e. ones that could trigger right-to-left Laxing Harmony). Following the literature, we expect that Laxing Harmony applies right-to-left and that the community-level grammar will show evidence of both Local Iterative Harmony and Non-Local Harmony. The factors for potential conditioning vowels informs our inference of Laxing Harmony processes as follows:

- The factor for the vowel in the initial syllable suggests left-to-right Local Iterative Harmony in which neutral vowels are transparent. This is because laxing in the initial syllable (e.g. due to Closed-Syllable Laxing or due to Non-Local Harmony) predicts tense/lax quality regardless of intervening vowels. For example, the /y/ in /vylnerabilite/ *vulnérabilité* 'vulnerability' could lax due to optional Closed-Syllable Laxing and trigger the laxing of /i/ in both later syllables despite the intervening mid-high /e/ and low /a/, yielding [vʏlnerabɪlɪte].
- The factor for preceding vowels suggests left-to-right Local Harmony, which is Iterative if all models (i.e. all syllables) show the effect. Neutral vowels are opaque unless the presence of a lax high vowel in the initial syllable also has a significant laxing effect. For example, the

/i/ in /ʒyrisprydɑ̃s/ *jurisprudence* 'jurisprudence' could undergo optional Closed-Syllable Laxing and trigger laxing of the following syllable's /y/, yielding [ʒyrɪsprʏdɑ̃s].

- The factor for following vowels suggests right-to-left Local Harmony. If all models (not only the penult models) show the effect, then Local Harmony is likely to be Iterative; if the model for only a single syllable (e.g. penults) shows this effect, then Local Harmony is likely to be Non-Iterative. Neutral vowels are opaque unless the presence of a lax high vowel in the final syllable also has a significant laxing effect. For example, the /y/ in the second syllable of /yzyrpasjɔ̃/ *usurpation* 'usurpation' could undergo optional Closed-Syllable Laxing and trigger laxing in the preceding syllable's /y/, yielding [ʏzʏrpasjɔ̃]. If there is a significant effect of the vowel in the following syllable only for models analysing the penultimate syllable, then /ilymin/ *illumine* 'illuminates' is likely to be realised as [ilʏmɪn] if Local Harmony is not Iterative; if the effect is found across models, then [ɪlʏmɪn] is expected if Local Harmony is Iterative.
- The factor for the vowel in the final syllable suggests right-to-left Local Iterative Harmony with transparent neutral vowels if the effect is found in all models, but right-to-left Non-Local Harmony if the effect is found only in initial syllables. For example, the /i/ in /akyzatif/ *accusatif* 'accusative' could lax as a result of obligatory Closed-Syllable Laxing and trigger laxing in /y/ due to Local Iterative Harmony, yielding [akʏzatɪf]. If the final syllable has a significant effect only in initial syllables and the following syllable does not have a significant effect, then we expect that /ilymin/ *illumine* 'illuminates' is likely to be realised as [ɪlymɪn] (Non-Local Harmony), just like /inedit/ *inédite* 'unpublished' is likely to be realised as [ɪnedɪt].

We are interested not only in whether Laxing Harmony applies, but also in the possibility that laxing results from (gradient) coarticulation. This question motivates our decision to include vowel height and not solely whether potential conditioning vowels are (a) tense high vowels, (b) lax high vowels, or (c) neutral (non-high) vowels. We expect that coarticulation partly conditions laxing and therefore that a high vowel is more likely to be realised as lax the lower the vowels in neighbouring syllables are. For example, if laxing is in part the result of coarticulation, laxing should be more likely to occur in /vitɛs/ *vitesse* 'speed' (mid-low vowel) than in /vito/ *vitaux* 'vital (pl.)' (mid-high vowel), and more likely to occur in /vital/ *vital* 'vital' (low vowel) than in either of those words.

Finally, for disyllabic words, we consider the possibility of Laxing Disharmony, as in the case of words like /midi/ *midi* 'noon', which yields [mɪdi]. Laxing Disharmony is said to be conditional upon the two vowels being featurally identical (underlyingly) and the final vowel not being laxed. Based on this, we include a factor that indicates whether the token is from a disyllabic word with two phonemically identical high vowels in open syllables. We expect that this context will favour laxed initial-syllable vowels due to the dissimilation phenomenon reported by Dumas (1987) and Poliquin (2006), meaning that initial-syllable laxing is more likely in a word like /midi/ *midi* 'noon' than /minu/ *minou* 'small cat'.

3.4.4.2 Syllable structure

Following Poliquin (2006), we consider the possibility that laxing may be triggered by syllable structure, rather than by harmony. For example, Poliquin suggests that laxing is possible in the initial syllables of /mistɛr/ *mystère* 'mystery' and /sulri/ *soulerie* 'drinking spree' because this syllable is closed. To test this, we wrote an R script that implemented an onset maximisation procedure for syllabification based on the word's surface form, allowing us to code each syllable as being open or closed. Following both the literature and initial data exploration, we only include this factor for the syllable under examination (i.e. we do not include the syllable structure of the penult as a factor in the model for the initial syllable of trisyllabic words).

We additionally test Cuerrier and Reiss's (2018) proposal that derived geminates are present in Laurentian French and that these geminates trigger optional Pre-Geminate Laxing in word-initial lax high vowels in words like /il(:)isit/ [ɪl(:)isɪt]~[il(:)isɪt] *illicite* 'illicit' and /il(:)egal/ [ɪl(:)egal]~ [il(:)egal] *illegal* 'illegal'. In order to test Pre-Geminate Laxing, we use orthographic cues to predict likely underlying geminates that are located at morpheme boundaries (and therefore most likely to be in speakers' representations). We code words for the presence of the historical *-in* prefix for negation, which consistently contains a high vowel and which seems especially likely to be associated with a geminate by speakers due to the morpheme's frequency. If a word begins with "i" followed by an orthographicall`y doubled "l", "m", "n" or "r" (the segments most often associated with the negation prefix), then it was coded as potentially having a geminate; otherwise, the initial syllable was strictly coded as being open or closed, based on its surface form. Following Cuerrier and Reiss, we expect that words with this prefix, like *illicite*,

will be more likely to undergo initial-syllable laxing than other words, thereby supporting Pre-Geminate Laxing.[38]

3.4.4.3 Lengthening Consonants

As described in section 3.1, high vowels in syllables closed by a Lengthening Consonant other than /r/ (i.e. /v z ʒ vr/) may undergo retensing or otherwise trigger Laxing Harmony even when they surface as tense. Lengthening triggered by Lengthening Consonants is attested in non-final syllables in Laurentian French, albeit potentially only in derived words (Walker, 1984). As such, we include as a factor for disyllabic words whether the medial consonant belongs to the class of Lengthening Consonants. We did this for disyllabic words only, as preliminary data exploration and initial models suggested it was not warranted for longer words. For instance, the penult in /vizit/ *visite* 'visit' was coded as being followed by an intervocalic Lengthening Consonant, suggesting the penult vowel could be lengthened and therefore more likely to be realised as tense. We expect that intervocalic Lengthening Consonants can trigger lengthening and therefore reduce laxing rates, making laxing in /vinil/ *vinyle* 'vinyl' more likely than in /vizit/ *visite* 'visit'.

3.4.4.4 Morphological structure

Many effects – for example that of Lengthening Consonants – could be sensitive to the morphological structure of a word. As a result, we code whether (a) the token is base-final, and whether (b) the base-final syllable could lax when there are no additional suffixes. For example, the penult in /visit+e/ *visiter* 'to visit' is coded as base-final. Furthermore, the initial syllable is coded as being in a word where it could have laxed (cf. /visit/ *visite* 'visit.1SG.PRS', where obligatory Closed-Syllable Laxing occurs in the final syllable and thereby could trigger harmony to its left). We expect that base-final syllables will be unlikely to be realised as lax due to Retensing from Resyllabification, which is expected to occur in most of these tokens.

[38] We do not directly test whether other indications confirm the presence of a geminate (e.g. differences in consonant duration, as in Cuerrier and Reiss, 2017); instead, we only consider whether laxing occurs at an increased rate before potential geminates.

In disyllabic words, we additionally expect that a base-final syllable is unlikely to be lax if it is followed by a Lengthening Consonant because Lengthening could have occurred at an earlier stage of the derivation. However, in words of all lengths, we predict that high vowels are more likely to be lax if the base-final vowel would have laxed in non-derived forms, consistent with Poliquin's (2006) observation that the retensed vowel can trigger harmony as if it were lax due to Laxing Harmony applying at an earlier stage of the derivation.

3.4.4.5 Speech rate

As a final predictor, we include speech rate. This was calculated based on the number of syllables per second, and then the speech rate was converted to z-scores individually for each speaker (to determine if the speech is fast for a given speaker, rather than just fast in general). Given that faster speech rates are typically associated with greater reduction (e.g. Spilková, 2014) and given that we expect that gradient vowel reduction influences laxing, we expect that faster speech rates will be associated with more frequent laxing. Furthermore, speech rate allows us to infer whether the pattern described as harmony is more likely to be gradient coarticulation or category-changing harmony. If speech rate increases a gradient effect of vowel height (laxing is favoured more adjacent to lower vowels), then that laxing is interpreted as being coarticulatory. We expect that laxing is both phonetic and phonological; we expect to find clear differences in effect size and direction whereby lax high vowels trigger laxing considerably more than other vowels, while also finding gradient vowel height effects that are magnified at faster speech rates.

3.4.5 Summary

In this study, we train a forced-aligner on categorical final-syllable laxing in order to use that forced-aligner to classify the tense/lax quality of high vowels in non-final syllables. We then perform mixed-effects logistic regression to determine which factors influence the tenseness of those high vowels in non-final syllables. We consider the effects of other vowels from the same word (Laxing Harmony, coarticulation), syllable structure (Closed-Syllable Laxing, potential geminates), Lengthening Consonants (Pre-Fricative Tensing, Lengthening Consonants), and

morphology (Retensing from Resyllabification, opaque Laxing Harmony) in our analysis. As we show in the following section, these factors allow us to better understand what conditions the laxing of high vowels outside of final syllables at the community level, and therefore what language learners likely receive as input when acquiring high-vowel laxing.

3.5 Results

As we previously noted, the interpretation of factors and the factors that can be tested vary by both the syllable number and the size of the word. We therefore organise the results by statistical model, as the models were designed around those differences. We first examine non-final syllables in disyllabic words (section 3.5.1), then trisyllabic words (section 3.5.2 for initial syllables and section 3.5.3 for penults), and finally words of four or more syllables (sections 3.5.4, 3.5.5 and 3.5.6 for initial syllables, medial syllables and penults, respectively). Finally, section 3.5.7 will provide a thematic overview of the overall effects of each factor. Table 11 provides a summary of the following model information: the token count, the type (unique word) count, the speaker count, model fit criteria (AIC and BIC), and variance summaries for random effects (words and speakers). We note that all models show more variation across words than across speakers, which provides suggestive evidence that Poliquin's (2006) categorisation of speakers into distinct Laxing Harmony groups may overstate the differences between speakers. We will also include the amount of data in each model at the start of the appropriate section to indicate the degree of confidence one should have in interpreting the model.

Table 11: Model details for each of our six models.

Model	Tokens (Types)	Speaker	AIC	BIC	Word variance (std. deviation)	Speaker variance (std. deviation)
initial syllable, 2 syllables (§3.5.1)	10,749 (1,334)	131	8780.5	8918.9	3.2607 (1.8057)	0.1692 (0.4114)
initial syllable, 3 syllables (§3.5.2)	2,140 (560)	124	2053.4	2189.4	2.861 (1.6914)	0.280 (0.5292)
penultimate syllable, 3 syllables (§3.5.3)	6,691 (1,350)	131	6329.3	6492.7	3.5985 (1.8969)	0.1702 (0.4126)
initial syllable, 4+ syllables (§3.5.4)	1,336 (210)	118	955.2	1069.6	23.405 (4.838)	1.876 (1.370)
medial syllables, 4+ syllables (§3.5.5)	2,484 (519)	125	2164.1	2344.4	7.043 (2.6539)	0.401 (0.6333)
penult syllables, 4+ syllables (§3.5.6)	2,083 (497)	127	1787.3	1962.1	13.5324 (3.6786)	0.6321 (0.7951)

3.5.1 Initial syllables of disyllabic words

Table 12 shows the model output for disyllabic words, which was based on 10,749 tokens (1,334 unique words) from 131 speakers. We use the models for disyllabic words to test Laxing Harmony and gradient height effects, syllable shape effects (Closed-Syllable Laxing and Retensing from Resyllabification), Pre-Fricative Tensing, Laxing Disharmony, and finally speech rate effects.

We begin with the question of category-changing phonological effects (Laxing Harmony) of potential conditioning vowels. We expected to find that lax high vowels in the final syllable would be associated with high vowels in the initial syllable laxing more often (§3.4.4.1) due to Laxing Harmony (§3.2.5). Counter to expectation, we observe that the initial vowel has a tendency to be tense even when the final syllable contains a laxed high vowel (this corresponds to the intercept in Table 7: β=-0.7185, p=0.0177). However, we do find that the penult is less likely to lax if the final-syllable vowel is neither high nor lax. This suggests that the presence of a laxed high vowel in the final syllable (e.g. as expected in /minim/ *minime* 'minimal', where the final syllable categorically laxes) does favour laxing in the initial syllable relative to there being a vowel of another quality in the final syllable (e.g. the /œ/ in /minœr/ *mineur* 'minor' or the /a/ in /minabl/ *minable* 'pathetic'). However, a high lax vowel is not the majority variant even in

words with the appropriate phonological shape for Laxing Harmony, so that a word like *minime* is more likely to be pronounced as [minɪm] than as [mɪnɪm].

Turning to syllable shape, we expected that non-final high vowels would lax more often in closed syllables (§3.4.4.2) as a result of optional Closed-Syllable Laxing (§3.2.3), while base-final syllables would be associated with lower laxing rates (§3.4.4.4) due to Retensing from Resyllabification (§3.2.5.5). However, neither syllable shape nor morphological structure significantly affects laxing rates in disyllabic words. Contrary to expectations, laxing is not significantly affected by syllable shape in disyllabic words, which means that the laxing rates in the initial open syllable in /misil/ *missile* 'missile' and the initial closed syllable in /mistik/ *mystique* 'mystical' are not significantly different with respect to their laxing rates. Syllable structure effects are not significantly affected by morphological structure, indicating that /mit+ik/ *mythique* 'mythical' does not exhibit significantly different laxing rates in the initial syllable compared to *missile* (or, consequently, *mystique*). Furthermore, inferring more generally from the lack of significant effect from morphological structure, words like /sit+e/ *cité* 'cited' and /site/ *cité* 'city' are not significantly different with respect to their laxing rates even though the former word would exhibit laxing in its base in non-derived contexts ([sɪt]) and would then undergo Retensing from Resyllabification.

Next we examine the factors that are only applicable for disyllabic words: lengthening consonants (except /r/) and Laxing Disharmony. With respect to lengthening consonants in intervocalic position, we expected to find that high vowels followed by a Lengthening Consonant as the only interconsonantal segments would be associated with lower laxing rates (§3.4.4.3) due to Pre-Fricative Tensing (§3.2.1). Consistent with this, we find the predicted preference for a tense vowel preceding such a consonant (β=-0.6509, p=0.0295). As a result, the initial syllable of /visit/ *visite* 'visit' is less likely to undergo laxing than the initial syllable of /vinil/ *vinyle* 'vinyl'; again, this pattern is not sensitive to morphological structure. With respect to words in which both syllables contain identical high-vowel phonemes in open syllables, like /midi/ *midi* 'noon', we expected to find increased laxing rates (§3.4.4.1) due to Laxing Disharmony (§3.2.4). However, we find no significant effect to counter the effect of a tense high vowel in the final syllable of these words, though there is a trend in the expected direction. The lack of significant result is not likely to be due to a small sample size: there were 1164 tokens (40 distinct words)

with the appropriate shape for dissimilation to potentially occur (albeit from a relatively small number of lemmas, given limitations of the lexicon).

Table 12: Model output for disyllabic words.

	Coefficient	Std. Error	t-value	p-value	
(Intercept)	-0.7185	0.3030	-2.371	0.0177	*
Closed syllable	-0.1681	0.3144	-0.535	0.5928	
Base-final	-0.0648	0.2217	-0.292	0.7701	
Faster speech	0.4092	0.0653	6.264	<0.0001	***
Lengthening Consonant	-0.6509	0.2990	-2.177	0.0295	*
Final tense high vowel	-1.2369	0.3322	-3.723	0.0002	***
Final mid-high vowel	-1.2743	0.2735	-4.660	<0.0001	***
Final mid-low vowel	-1.8008	0.2670	-6.744	<0.0001	***
Final low vowel	-1.0383	0.2894	-3.587	0.0003	***
Dissimilation context	0.8228	0.7046	1.168	0.2429	
Base-final, Dissimilation context	-0.1098	0.5854	-0.188	0.8512	
Base-final, Closed syllable	0.0192	0.3536	0.054	0.9566	
Base-final, Lengthening Consonant	0.0548	0.3132	0.175	0.8611	
Faster speech, Final tense high vowel	-0.3746	0.0898	-4.170	<0.0001	***
Faster speech, Final mid-high vowel	-0.2602	0.0906	-2.873	0.0041	**
Faster speech, Final mid-low vowel	-0.3047	0.0926	-3.290	0.0010	**
Faster speech, Final low vowel	-0.2633	0.1034	-2.545	0.0109	*

Finally, we considered factors that would suggest that laxing results from phonetic pressures. The first of these is that we expected to find gradient effects of vowel height from the potential conditioning vowel as a result of coarticulation (§3.4.4.1). Contrary to our expectation that high-vowel laxing would be sensitive to coarticulatory pressures such that lower potential conditioning vowels would be associated with more laxing (§3.4.4.1), the likelihood of laxing a high vowel in the initial syllable decreases as the potential conditioning vowel gets lower. Low vowels go against this trend, however; final-syllable low vowels are associated with higher likelihoods of being preceded by a laxed vowel compared to mid-high and mid-low vowels (although still disfavouring laxing relative to lax high vowels; β=-1.0383, p=0.0003). The second is that faster speech rates would be associated with higher rates of laxing due to phonetic reduction (§3.4.4.5) and that faster speech rates would additionally magnify gradient vowel height effects due to coarticulation (§3.4.4.5). Seemingly consistent with this, we find that faster speech rates are associated with a higher likelihood of harmonising with a following lax high

vowel (β=0.4092, p<0.0001), which could indicate phonetic pressure to lax. However, faster speech is also associated with a reduced likelihood of harmonising with vowels of other qualities with no indication of lower vowels conditioning high-vowel laxing most at faster speech rates (coefficients from -0.2602 to -0.3746, with the highest p-value at p=0.0109). This result suggests the possibility of a phonetic component to laxing (speech rate effect before high lax vowels), but we will argue in section 3.6 that this is not actually the case.

3.5.2 Initial syllables of trisyllabic words

Table 13 shows the model output for the initial syllable of trisyllabic words, based on 2,140 tokens (560 unique words) from 124 speakers. As will be the case for other models based on words of three or more syllables, we use our models of the initial syllable of trisyllabic words to test Laxing Harmony, syllable shape effects (Closed-Syllable Laxing and Pre-Geminate Laxing) and speech rate effects.

Beginning with factors capturing Laxing Harmony, we expected that lax high vowels in either non-initial syllable would be associated with higher laxing rates in the initial syllable (§3.4.4.1, §3.4.4.4) due to Laxing Harmony triggered either by the following vowel due to right-to-left Local Harmony or by the final-syllable vowel due to either Non-Local Harmony or Local Iterative Harmony (§3.2.5). However, in the initial syllable of trisyllabic words, we find no significant preference for high vowels to surface as lax when the other vowels in the word are high lax vowels (intercept: β=0.4590, p=0.3569), as shown in Table 8. Consistent with our expectation that Laxing Harmony motivates high-vowel laxing, we do observe significant decreases in the likelihood of an initial high vowel to surface as lax when either other vowel in the word is not a high lax vowel. This means that laxing is more likely in the initial syllable of /difisil/ *difficile* ‘difficult’ if the medial syllable also undergoes laxing (yielding [dzɪfisɪl]). Turning to morphologically complex words, the base-final vowel is not a significant predictor of tenseness. This result suggests that base-final vowels do not have special status compared to other vowels as triggers for Laxing Harmony. In other words, laxing in the penult increases the likelihood of laxing in the initial syllable, but /myzik+al/ *musical* ‘musical’ is unlikely to be realised with lax [ʏ] in the initial syllable unless the penult vowel is realised as lax [ɪ] (contrary to expectations from Retensing from Resyllabification based upon which the vowel-initial suffix

would motivate resyllabification of a base-final coda and therefore also retensing of a high vowel at the end of the base (§3.2.5.5)).

Turning to syllable structure effects, we expected that closed syllables would be associated with higher laxing rates (§3.4.4.2) due to Closed-Syllable Laxing (§3.2.3), and that high vowels followed by underlying geminates would surface as lax more often (§3.4.4.2) due to Pre-Geminate Laxing (§3.2.5.3). On one hand, the surface syllable shape of the initial syllable has no significant effect on its own (i.e. the syllable being open or closed is not a significant predictor). On the other hand, the presence of an initial geminate sequence does significantly and dramatically increase the likelihood of laxing the initial high vowel (β=1.7055, p=0.0179). As such, the initial syllable of /il(:)isit/ *illicite* 'illicit', where the /l/ is a candidate for being a derived geminate, is more likely to contain a lax vowel than the initial syllable of /difisil/ *difficile* 'difficult'.

Table 13: Model output for the initial syllable of trisyllabic words.

	Coefficient	Std. Error	t-value	p-value	
(Intercept)	0.4590	0.4982	0.921	0.3569	
Faster speech	0.0899	0.2989	0.301	0.7637	
Initial geminate	1.7055	0.7200	2.369	0.0179	*
Open syllable	-0.2489	0.2801	-0.888	0.3743	
Tense high penult vowel	-0.9390	0.3230	-2.907	0.0037	**
Tense mid penult vowel	-1.2305	0.4012	-3.067	0.0022	**
Lax mid penult vowel	-1.4078	0.4581	-3.073	0.0021	**
Low penult vowel	-0.8336	0.4049	-2.059	0.0395	*
Tense high final vowel	-1.7281	0.8282	-2.087	0.0369	*
Tense mid final vowel	-2.3022	0.7605	-3.027	0.0025	**
Lax mid final vowel	-2.3753	0.8483	-2.800	0.0051	**
Low final vowel	-2.2960	0.7671	-2.993	0.0028	**
Base-final vowel could lax	0.8845	0.7961	1.111	0.2666	
Faster speech, Tense high penult vowel	0.1500	0.2541	0.591	0.5548	
Faster speech, Tense mid penult vowel	-0.0322	0.2859	-0.113	0.9104	
Faster speech, Lax mid penult vowel	0.0537	0.2841	0.189	0.8500	
Faster speech, Low penult vowel	-0.0728	0.2832	-0.257	0.7972	
Faster speech, Tense high final vowel	0.0677	0.3123	0.217	0.8284	
Faster speech, Tense mid final vowel	0.1230	0.2699	0.456	0.6486	
Faster speech, Lax mid final vowel	-0.0812	0.2472	-0.328	0.7426	
Faster speech, Low final vowel	-0.1414	0.2881	-0.491	0.6236	

Finally, we expected to find a gradient effect whereby a potential conditioning vowel is associated with higher laxing rates the lower that vowel is due to gradient coarticulatory effects (§3.4.4.1) and that this effect would be magnified at faster speech rates (§3.4.4.5). The coefficients combined with the standard error for the final-syllable vowel reveal that a final tense high vowel disfavours laxing the least relative to other final-syllable vowels (β=-1.7281, p=0.0369; cf. the other vowels at a minimum β=-2.3022 with similar standard errors). For the penult vowel, however, lower vowels increasingly lower the likelihood of laxing an initial high vowel, with the exception of low vowels that show an effect similar to high vowels. For example, /inɛgzakt/ *inexact* 'inexact' may be slightly more likely to have a word-initial lax vowel than /inegal/ *inégal* 'unequal' due to the quality of the word-medial vowel after accounting for other explanations for laxing rates. While this suggests a gradient phonetic effect, faster speech rates neither significantly affect laxing on their own (contrary to expectations if laxing were the result of phonetic reduction; §3.4.4.5), nor do they interact with penult or final vowels to suggest that there are strong coarticulatory pressures (§3.4.4.5).

3.5.3 Penults of trisyllabic words

Table 14 shows the model output for penults of trisyllabic words, based on 6,691 tokens (1,350 unique words) from 131 speakers. We use our models of the penult of trisyllabic words to test potential conditioning vowels (Laxing Harmony, coarticulation), syllable shape effects (Closed-Syllable Laxing) and speech rate effects.

We begin with the question of category-changing (Laxing Harmony) and gradient (coarticulation) effects of potential conditioning vowels, expecting to find that potential conditioning vowels in the final syllable will predict laxing rates (§3.4.4.1) due to right-to-left Laxing Harmony (§3.2.5). We additionally expect that the potential conditioning vowel in the initial syllable will influence laxing rates, but that this will mainly be due to gradient coarticulatory pressures rather than to the presence of left-to-right Laxing Harmony (§3.4.4.1). Looking at the intercept to infer the likelihood of laxing a medial vowel with lax high vowels in both adjacent syllables (e.g. [ɪlɪsɪt] or [ɪlisɪt] for /ilisit/ *illicite* 'illicit'), there is no significant preference to lax or not when the high vowel is flanked by syllables containing high lax vowels (β=-0.0428, p=0.9169). Regarding effects of adjacent vowels, we find that the strongest effects

are those of the initial syllable's vowel rather than of the final syllable's vowel, contrary to expectations. High vowels in the penult are more likely to lax if the initial vowel is a high lax vowel (e.g. due to Non-Local Laxing Harmony or Closed-Syllable Laxing). For example, the /y/ in /buskylad/ *bousculade* 'hustle' is more likely to undergo laxing if /u/ in the initial syllable does, while laxing in the final syllable of /bazilik/ *basilic* 'basil' does not significantly condition the first /i/ to lax. This runs counter to our expectations: we expected that initial vowels would mainly be associated with coarticulatory effects due to the right-to-left direction of Laxing Harmony (§3.4.4.1), but in trisyllabic words, conditioning vowels on the left – but not on the right – are the best predictors of vowel-conditioned laxing.

Table 14: Model output for the penultimate syllable of trisyllabic words.

	Coefficient	Std. Deviation	t-value	p-value	
(Intercept)	-0.0428	0.4101	-0.104	0.9169	
Faster speech	-0.4738	0.3047	-1.555	0.1199	
Closed syllable	0.5569	0.2082	2.674	0.0075	**
Tense high initial vowel	-0.8731	0.2957	-2.953	0.0031	**
Tense mid initial vowel	-0.6551	0.3290	-1.991	0.0465	*
Lax mid initial vowel	-0.7543	0.3483	-2.166	0.0303	*
Low initial vowel	-0.4481	0.3276	-1.368	0.1713	
Tense high final vowel	-1.4640	1.5900	-0.921	0.3572	
Tense mid final vowel	-2.5021	1.5721	-1.592	0.1115	
Lax mid final vowel	-2.0076	1.5943	-1.259	0.2080	
Low final vowel	-2.2759	1.5777	-1.443	0.1492	
Base-final vowel could lax	0.9098	1.5930	0.571	0.5679	
Vowel is base-final	1.0938	1.5873	0.689	0.4908	
Faster speech, Tense high initial vowel	0.5219	0.3152	1.656	0.0978	.
Faster speech, Tense mid initial vowel	0.6869	0.2870	2.393	0.0167	*
Faster speech, Lax mid initial vowel	0.6806	0.2978	2.285	0.0223	*
Faster speech, Low initial vowel	0.5261	0.2864	1.837	0.0662	.
Faster speech, Tense high final vowel	-0.0350	0.1826	-0.192	0.8479	
Faster speech, Tense mid final vowel	0.1502	0.1617	0.929	0.3531	
Faster speech, Lax mid final vowel	-0.1503	0.1751	-0.858	0.3906	
Faster speech, Low final vowel	-0.0622	0.1807	-0.344	0.7308	

Turning to other factors that could predict high-vowel laxing, we consider both syllable shape effects and speech rate. Regarding syllable shape, we expected to find that closed syllables favoured higher laxing rates (§3.4.4.1) due to Closed-Syllable Laxing (§3.2.3). Consistent with

this expectation, high vowels in the penult are more likely to surface as lax when that syllable is closed (β=0.5569, p=0.0075), indicating that /y/ is less likely to undergo laxing in /kɔmynjɔ̃/ *communion* 'communion' than in /kɔ̃bystjɔ̃/ *combustion* 'combustion'. Paired with this expectation was that base-final syllables would be associated with lower laxing rates (§3.4.4.4) due to Retensing from Resyllabification (§3.2.5.5). However, we do not find a significant effect associated with the morphological status of the word's medial syllable. Regarding speech rate effects, we expected that faster speech rates would be associated with higher laxing rates due to phonetic reduction and increased coarticulation (§3.4.4.5). The effect of potential conditioning vowels in the initial syllable does appear to be mediated by speech rate (though speech rate is only marginal as a main effect); in faster speech, laxing becomes more frequent after non-high vowels, indicating that /i/ in /fofile/ *faufiler* 'sneak out' is more likely to undergo laxing at faster speech rates under the influence of /o/ in the preceding syllable. Final vowels, on the other hand, have no significant effects, though there is a trend for lower rates of laxing when the final vowel is not a high lax vowel.

3.5.4 Initial syllables of longer words

Table 15 shows the model output for the initial syllable of words with at least four syllables, based on 1,336 tokens (210 unique words) and 118 speakers. We use our models of the initial syllable of words with four or more syllables to test syllable shape effects (Closed-Syllable Laxing), Laxing Harmony and speech rate effects.

We begin with the question of category-changing (Laxing Harmony). First, we expected that high vowels would be more likely to surface as lax when the following syllable or the final syllable contained a high lax vowel (§3.4.4.1) due to Laxing Harmony (§3.2.5). In the current model, however, we find only limited evidence from the potential conditioning vowel in the following syllable, and final-syllable vowel height is not a significant predictor of laxing rates. As such, we can infer that the initial syllable is significantly less likely to surface as lax in /nymerotasjɔ̃/ *numérotation* 'numbering' than in /simyltanemɑ̃/ *simultanément* 'simultaneously' (realised with laxing in the second syllable), but we cannot conclude that laxing in the second syllable in *simultanément* affects laxing in the initial syllable. However, given the limited data, we cannot strongly refute our expectations that a high lax vowel in the following syllable would

trigger laxing through Laxing Harmony: we find significant laxing effects for mid vowels in the following syllable, which is consistent with a combination of gradient height effects (in which low vowels do not participate) combined with a weaker trend to lax when the following syllable contains a high vowel. As such, it may be that our result stems from a combination of phonetic coarticulation and right-to-left Local Harmony.

Table 15: Model output for the initial syllable of words with four or more syllables.

	Coefficient	Std. Deviation	t-value	p-value	
(Intercept)	-2.2755	2.0985	-1.084	0.2782	
Faster speech	-1.8019	1.8953	-0.951	0.3417	
Closed syllable	2.0722	0.8624	2.403	0.0163	*
High tense next vowel	-0.4511	0.6189	-0.729	0.4660	
Mid tense next vowel	-2.2004	1.1143	-1.975	0.0483	*
Mid lax next vowel	-3.5632	0.8587	-4.150	0.0000	***
Low next vowel	1.6723	1.2899	1.296	0.1948	
High tense final vowel	-2.1526	2.1747	-0.990	0.3223	
Mid tense final vowel	-1.3703	1.9623	-0.698	0.4850	
Mid lax final vowel	-2.7464	2.1999	-1.248	0.2119	
Low final vowel	-0.9682	2.0289	-0.477	0.6332	
Base-final vowel could lax	1.9098	1.0995	1.737	0.0824	.
Faster speech, High tense next vowel	-0.2013	0.6183	-0.326	0.7448	
Faster speech, Mid tense next vowel	-0.3602	0.7062	-0.510	0.6100	
Faster speech, Mid lax next vowel	1.3515	0.8768	1.541	0.1232	
Faster speech, Low next vowel	0.0953	0.6556	0.145	0.8845	
Faster speech, High tense final vowel	1.1348	1.8282	0.621	0.5348	
Faster speech, Mid tense final vowel	1.7393	1.8422	0.944	0.3451	
Faster speech, Mid lax final vowel	2.1410	1.8509	1.157	0.2474	
Faster speech, Low final vowel	2.0750	1.9000	1.092	0.2748	

Second, we expected that the potential for a base-final syllable to contain a lax high vowel at any stage of the derivation would similarly be associated with increased laxing rates (§3.4.4.4) due to Laxing Harmony applying even when the result is opaque surface forms (§3.2.5.4; §3.2.5.5). We observe a marginal result of a potential conditioning vowel in the base-final syllable that is consistent with this expectation (β=1.9098, p=0.0824). For example, the final /i/ in /imaʒin/ *imagine* 'imagine' would undergo Closed-Syllable Laxing and can therefore trigger laxing of the initial /i/ regardless of whether that final-syllable /i/ undergoes Retensing from

Resyllabification, increasing the likelihood of pronouncing /imaʒin+e/ *imaginer* 'to imagine' as [ɪmaʒine] rather than as [imaʒine].

Turning to syllable structure effects, we expected to find that closed syllables are associated with significantly higher laxing rates (§3.4.4.2) due to Closed-Syllable Laxing (§3.2.3). We do again find a significant effect for syllable structure, with closed syllables being associated with higher rates of laxing (β=2.0722, p=0.0163). We can therefore infer that /y/ is more likely to undergo laxing in /misteriøzmɑ̃/ *mystérieusement* 'mysteriously' than in /sinematograf/ *cinématographe* 'cinematographer'. This model was our smallest in terms of token count and type count and therefore we cannot confidently infer the absence of processes from null results for other factors.

Finally, we consider speech rate effects. First, we expected that high vowels would surface as lax more often at faster speech rates due to phonetic reduction (§3.4.4.5). However, speech rate has no significant main effect. Second, we expected that speech rate would mediate coarticulatory effects such that faster speech would magnify gradient coarticulatory effects. Once again, no significant interaction was found; we thus cannot conclude that coarticulation results in higher laxing rates at faster speech.

3.5.5 Medial syllables of longer words

Table 16 shows the model output for the initial syllable of words with at least four syllables, based on 2,484 tokens (519 unique words) and 125 speakers. We use our models of the medial syllable of words with four or more syllables to test potential conditioning vowel effects (Laxing Harmony, gradient coarticulation), syllable shape effects (Closed-Syllable Laxing) and speech rate effects (phonetic reduction, gradient coarticulation).

We begin with the effects of potential conditioning vowels. First, we expected to find that lax high vowels later in the word – i.e. in the following and final syllables – would be associated with increased laxing rates (§3.4.4.1) due to right-to-left Laxing Harmony (§3.2.5). In line with these expectations, we do not find a consistently significant effect of the previous syllable's vowel, suggesting that the laxing of /i/ in /kylminasjɔ̃/ *culmination* 'culmination' is not significantly associated with whether /y/ laxes in the initial syllable. However, we also do not

find a consistently significant effect of the following syllable's vowel. Instead, only the final syllable is associated with the expected categorical effect whereby lax high vowels in a potential conditioning syllable are associated with substantially higher rates of laxing than other vowels (marginal for high vowels and significant for all other vowel heights). As such, we tentatively infer that /i/ in the second syllable is more likely to undergo laxing in /meritokratik/ *méritocratique* 'meritocratic' than in /mediterane/ *méditerranée* '(the) Mediterranean', and probably also more likely than in /meritokrasi/ *méritocracie* 'meritocracy'. The presence of a final vowel that is not both high and lax is associated with significantly lower laxing rates, though the effect is smaller when the final vowel shares either the height (marginal; β=1.9097, p=0.0922) or lax quality (β=-0.19329, p=0.0027) with that conditioning vowel.

We now turn to syllable structure. First, we expected that high vowels would be more likely to surface as lax in closed syllables (§3.4.4.2) due to Closed-Syllable Laxing (§3.2.3). However, we observe that the syllable being closed does not significantly predict laxing, although the trend is in the expected direction (β=0.5217, p=0.1566). As such, we cannot infer a difference in laxing rates for /i/ between /gravitasjɔnɛl/ *gravitationnel* 'gravitational' and /maʒistralmɑ̃/ *magistralement* 'masterfully'. Second, we expected that base-final syllables would be more likely to surface as tense (§3.4.4.4) due to Retensing from Resyllabification (§3.2.5.5). Consistent with this expectation, base-final syllables have significantly lower rates of laxing (β=-1.5815, p=0.0080). For example, the second /i/ in /visit+asjõ/ *visitation* 'visitation' is unlikely to surface as lax.

Finally, we turn to potential indication of phonetic motivations for laxing. First, we expected to find a gradient effect of vowel height due to coarticulation, such that high vowels surface as lax more often when potential conditioning vowels are lower (§3.4.4.1). While the effects of potential conditioning vowels are not consistently significant, our results for adjacent syllables suggest that potential conditioning vowels are increasingly associated with lax high vowels the lower the potential conditioning vowel is. For both preceding and following syllables, low vowels significantly increase the likelihood of medial-syllable high vowels surfacing as lax. Second, we expected that high vowels would be more likely to lax during faster speech due to phonetic reduction (§3.4.4.5). As in most previous models, however, we do not find a significant main effect of speech rate. Third, we expected that the gradient effect of the height of potential

conditioning vowels would be mediated by speech rate as a result of speech rate influencing coarticulation (§3.4.4.5). Contrary to expectations, we find no significant interactions between speech rate and potential conditioning vowels.

Table 16: Model output for the non-penultimate medial syllable of words with four or more syllables.

	Coefficient	Std. Deviation	t-value	p-value	
(Intercept)	1.0040	1.0367	0.968	0.3328	
Faster speech	-0.0119	1.1564	-0.010	0.9918	
Closed syllable	0.5217	0.3682	1.417	0.1566	
Base-final syllable could lax	0.1321	0.3375	0.391	0.6956	
Base-final syllable	-1.5815	0.5960	-2.653	0.0080	**
High tense previous vowel	-0.1564	0.3356	-0.466	0.6411	
Mid tense previous vowel	0.4634	0.3799	1.220	0.2225	
Mid lax previous vowel	-0.9806	0.5509	-1.780	0.0751	.
Low previous vowel	-1.0686	0.4095	-2.610	0.0091	**
High tense next vowel	-0.4765	0.3098	-1.538	0.1240	
Mid tense next vowel	-0.2534	0.5559	-0.456	0.6485	
Mid lax next vowel	0.3375	0.6501	0.519	0.6037	
Low next vowel	-1.3293	0.3749	-3.545	0.0004	***
High tense final vowel	-1.9097	1.1340	-1.684	0.0922	.
Mid tense final vowel	-3.0331	1.0097	-3.004	0.0027	**
Mid lax final vowel	-1.9329	0.9646	-2.004	0.0451	*
Low final vowel	-3.0600	1.0201	-3.000	0.0027	**
Faster speech, High tense previous vowel	-0.3295	0.3703	-0.890	0.3735	
Faster speech, Mid tense previous vowel	0.0646	0.3270	0.197	0.8435	
Faster speech, Mid lax previous vowel	0.0361	0.4219	0.086	0.9318	
Faster speech, Low previous vowel	-0.2103	0.3184	-0.661	0.5088	
Faster speech, High tense next vowel	0.3690	0.3373	1.094	0.2740	
Faster speech, Mid tense next vowel	-0.3937	0.4384	-0.898	0.3692	
Faster speech, Mid lax next vowel	-0.1681	0.3879	-0.433	0.6648	
Faster speech, Low next vowel	0.2069	0.3427	0.604	0.5461	
Faster speech, High tense final vowel	0.5003	1.1257	0.444	0.6567	
Faster speech, Mid tense final vowel	-0.2181	1.0950	-0.199	0.8421	
Faster speech, Mid lax final vowel	0.0027	1.0744	0.003	0.9980	
Faster speech, Low final vowel	-0.0247	1.0955	-0.023	0.9820	

3.5.6 Penults of longer words

Table 17 shows the model output for the initial syllable of words with at least four syllables, based on 2,083 tokens (497 unique words) and 127 speakers. We use our models of the penult in words with four or more syllables to test potential conditioning vowel effects (Laxing Harmony, gradient coarticulation), syllable shape effects (Closed-Syllable Laxing) and speech rate effects (phonetic reduction, gradient coarticulation). Overall, we observe that laxing in penults is particularly unlikely even in favouring contexts (i.e. with lax high vowels in all other syllables), as we can infer from the intercept (β=-4.6835, p=0.0023).

We begin with the effects of potential conditioning vowels. First, we expected that penultimate high vowels would be more likely to surface as lax when potential conditioning vowels are lax high vowels due to Laxing Harmony (§3.4.4.1; §3.2.5). However, perhaps due to the relatively small number of tokens relative to the number of predictors, we find no consistent significant effects to support this expectation of Laxing Harmony effects. Second, we also expected that potential conditioning vowels would be associated with gradient coarticulatory effects whereby lower vowels are more associated with high vowels surfacing as lax than higher conditioning vowels are (§3.4.4.1). Again, our results are inconclusive, though we do note that some trends suggest gradient height effects are present and that the two significant effects of the conditioning vowel are for mid-low and low conditioning vowels. Finally, we expected that the potential to lax in the base-final syllable would be associated with higher laxing rates due to Laxing Harmony applying at stages of the derivation when Closed-Syllable Laxing had already applied but neither Retensing from Resyllabification nor Pre-Fricative Tensing had yet (§3.4.4.4). Consistent with this expectation, we find that base-final syllables that could undergo laxing significantly increase the likelihood of laxing (β=1.6804, p=0.0048). It is worth noting that, in these cases, the base-final vowel often occurs to the left of the penult vowel, as in /pɔsibil+ite/ *possibilité* ‘possibility’, /elektris+ite/ *électricité* ‘electricity’, and /pyblis+ite/ *publicité* ‘publicity’, and therefore the target is also in a suffix (most often the first syllable of *–ité* ‘has the quality or characteristic of’).

Table 17: Model output for the penultimate syllable of words with four or more syllables.

	Coefficient	Std. Deviation	t-value	p-value	
(Intercept)	-4.6835	1.5392	-3.043	0.0023	**
Faster speech	-0.9873	1.9691	-0.501	0.6161	
Closed syllable	-0.1842	0.5482	-0.336	0.7369	
Base-final vowel could lax	1.6804	0.5956	2.821	0.0048	**
Base-final vowel	-0.1013	0.5039	-0.201	0.8407	
High tense initial vowel	0.1470	0.6484	0.227	0.8207	
Mid tense initial vowel	-0.1774	0.7874	-0.225	0.8218	
Mid lax initial vowel	0.1066	0.8155	0.131	0.8960	
Low initial vowel	1.4006	0.8079	1.734	0.0830	.
High tense previous vowel	-0.5551	0.4121	-1.347	0.1780	
Mid tense previous vowel	0.7111	0.6446	1.103	0.2700	
Mid lax previous vowel	1.8793	0.7611	2.469	0.0135	*
Low previous vowel	1.5244	0.6107	2.496	0.0126	*
High tense final vowel	-1.3319	1.5781	-0.844	0.3987	
Mid tense final vowel	-0.7659	1.2044	-0.636	0.5248	
Mid lax final vowel	-0.3220	1.2043	-0.267	0.7892	
Low final vowel	0.9281	1.2640	0.734	0.4628	
Faster speech, High tense initial vowel	0.0367	0.6985	0.053	0.9581	
Faster speech, Mid tense initial vowel	0.3723	0.6158	0.605	0.5454	
Faster speech, Mid lax initial vowel	0.0061	0.6622	0.009	0.9927	
Faster speech, Low initial vowel	0.3340	0.6414	0.521	0.6026	
Faster speech, High tense previous vowel	0.9700	0.4816	2.014	0.0440	*
Faster speech, Mid tense previous vowel	0.5196	0.4832	1.075	0.2823	
Faster speech, Mid lax previous vowel	-0.0264	0.5340	-0.049	0.9605	
Faster speech, Low previous vowel	0.4740	0.4743	0.999	0.3176	
Faster speech, High tense final vowel	-0.1163	2.0220	-0.058	0.9541	
Faster speech, Mid tense final vowel	0.4065	1.8546	0.219	0.8265	
Faster speech, Mid lax final vowel	0.7750	1.8552	0.418	0.6762	
Faster speech, Low final vowel	0.6325	1.8629	0.340	0.7342	

We now turn to syllable shape effects. First, we expected that closed syllables would be associated with higher laxing rates (§3.4.4.2) due to Closed-Syllable Laxing (§3.2.3). Contrary to this expectation, we do not find a significant effect of syllable structure for the penult of words with four or more syllables. As a result, we cannot conclude that the laxing of /i/ occurs at different rates in /kɔ̃tradiktwɑr/ *contradictoire* ‘contradictory’ compared to /kɔ̃pozitœr/ *compositeur* ‘composer’. Second, we expected that high vowels would be likely to surface as

tense in base-final syllables (§3.4.4.4) due to Retensing from Resyllabification (§3.2.5.5), like in /parazit+ɛr/ *parasitaire* 'parasitic'. However, we find no significant effect of syllables being base-final.

Finally, we consider speech rate effects. First, we expected that laxing would be more likely at faster speech rates due to phonetic reduction (§3.4.4.5). However, no significant main effect is found for speech rate. Second, we expected that speech rate would modulate coarticulatory effects such that gradient phonetic effects would be magnified at faster speech rates (§3.4.4.5). There is one significant interaction with speech rate, but we suspect this result to be spurious given the lack of systematic effects across contexts and for other interactions with speech rate.

3.5.7 Summary

In summary, we find recurring patterns with respect to the conditioning factors for laxing in non-final syllables. The most consistent of these patterns is that lax high vowels in syllables later in the word are associated with laxing, consistent with right-to-left Laxing Harmony proposed in the literature. We see this most evidently through the frequent effect of potential conditioning vowels in the following syllable, which suggests Laxing Harmony. Closed syllables do not consistently favour laxing across models for non-final syllables, which suggests a much weaker effect of Closed-Syllable Laxing than we had expected (though see section 3.6.1.1 for discussion). The effects of morphological structure are similarly varied in our models: while we find evidence for high vowels surfacing as tense in base-final syllables (suggestive of Retensing from Resyllabification) and for high vowels to lax under the influence of laxing having been available in the base-final syllable (suggestive of Laxing Harmony applying at earlier stages of the derivation), the empirical support is limited. Finally, speech rate effects are only rarely significant, both as a main effect and in interaction with the vowel quality of a neighbouring syllable's vowels.

Like previous work, the current study strives to elucidate the community-level grammar of category-changing phonological processes that affect the tense/lax quality of high vowels in Laurentian French. However, one observation that cannot be ignored is that lax high vowels can surface with no apparent trigger. For example, a substantial number of high vowels underwent laxing in open initial syllables with no high vowel elsewhere in the word. The next section will

therefore discuss the implications of our results, including the question of whether laxing is predominantly the result of phonological processes or instead results primarily from phonetic pressures, by highlighting the factors relevant to high-vowel laxing in Laurentian French and the implications of those factors.

3.6 Discussion

Overall, we find considerable variability in the tense/lax quality of high vowels with about 23% of high vowels in non-final syllables being lax, and not all tokens have clear triggers for laxing based on phonological processes proposed in the literature. In this section, we return to our two research questions. We address our first question in section 3.6.1, where we discuss which phonological processes a learner is expected to generalise based on the community-level input. We then turn to our second question in section 3.6.2, where we discuss the possibility that laxing results from phonetic pressures (reduction, coarticulation) rather than phonological ones (e.g. Laxing Harmony).

3.6.1 A Phonological Grammar of High-Vowel Tenseness

Many patterns emerge from a community-level examination of high-vowel tenseness in Laurentian French. This allows us to address our first research question, reiterated here:

> Which phonological processes determining high vowel tenseness can learners be expected to generate based on naturalistic input from the speech community?

In particular, we can directly address whether a learner would be expected to infer the application of phonological processes proposed in the literature: syllable shape effects like optional Closed-Syllable Laxing and geminate-triggered laxing (section 3.6.1.1), Laxing Disharmony (section 3.6.1.2) and Laxing Harmony (section 3.6.1.3). We end our discussion of the community-level phonological grammar with an examination of the effects of Lengthening Consonants (section 3.6.1.4).

3.6.1.1 Syllable shape effects

Overall, we find only limited evidence for non-final closed syllables favouring laxing, as this effect is less consistent across models than expected. There are three potential explanations. The first possibility is that the process being optional leads to insufficient data to infer Closed-Syllable Laxing, and therefore the model does not show a significant effect of laxing in closed syllables. This seems unlikely; closed syllables are common in the data and often words with closed syllables have no other trigger for laxing that the models could misinterpret laxing as originating from. For example, the initial syllable of /mistɛr/ *mystère* 'mystery' could not undergo laxing due to Laxing Harmony or Laxing Disharmony. The second possibility is that laxing is too frequent in open syllables to conclude that closed syllables have a significant laxing effect. This also seems unlikely; given that we included factors that should account for laxing in open syllables, Closed-Syllable Laxing should still have been detected. The third and final possibility is that our coding decisions for syllabification obfuscated the application of Closed-Syllable Laxing in non-final syllables. We argue in favour of this last interpretation, in particular because no opposite-direction effects (i.e. Closed-Syllable Tensing) were found, though we expect the frequency of laxing in open syllables may have further confounded the issue.

In our coding scheme, we followed Poliquin (2006) in syllabifying surface forms without distinguishing word-internal codas from onsets of syllables with a deleted nucleus. As such, our coding for closed syllables combined (a) consistently closed syllables, like the first syllable in /mistik/ [mɪs.tsɪk] *mystique* 'mystical', and (b) syllables closed as a result of vowel deletion, like the first syllable in /difisil/ *difficile* 'difficult' realised with High-Vowel Deletion, as in [dʑɪf.sɪl]. If Closed-Syllable Laxing only applies in closed syllables that do not result from the deletion of a following syllable's vowel, then our mixed results are partly expected – and the true Closed-Syllable Laxing effect could be stronger than what is suggested by this analysis. Based on the literature and on the consistently positive coefficients we find for laxing in closed syllables, we suggest that Closed-Syllable Laxing is likely to apply, albeit optionally, in non-final syllables. We leave for future work a comparison of the two types of closed syllables we identify here (closed as a result of High-Vowel Deletion or closed prior to High-Vowel Deletion), which may be best analysed separately for each speaker in order to determine whether individual speakers vary with respect to the type(s) of closed syllables that can trigger Closed-Syllable Laxing.

We now turn to Cuerrier and Reiss's (2018) proposal that geminates trigger laxing. For example, the initial vowels in /il(:)isit/ *illicite* 'illicit' and /il(:)egal/ *illegal* 'illegal' may be more likely to undergo laxing because the initial syllable may be closed by a geminate in Laurentian French. Based on the initial syllable of trisyllabic words – the only context where the hypothesis could reliably be tested – we find that geminates do favour laxing of the preceding vowel. The preference to lax in this context is, in fact, significantly greater than the preference to lax simply due to the syllable being closed. However, we also find that the existence of Non-Local Harmony is supported (thereby motivating laxing, see section 3.6.1.3) and our model for the initial syllable of trisyllabic words confirms these effects. This conclusion implies that our support for Cuerrier and Reiss's proposal is only partial, namely, their proposal that geminates provide an *alternative* explanation for Non-Local harmony. There is a potential confound in the distribution of probable geminates, however: while the presence of a morphological boundary increases our confidence that speakers could infer the presence of a geminate, we cannot confidently conclude that the additional effect is from the geminate rather than an effect associated with the morphological boundary itself. Furthermore, as mentioned in footnote 13, we do not test whether there is phonetic evidence of geminates in our data (e.g. longer consonant durations). Despite these two reservations, we believe it is likely that geminates have an effect at least equal to the effect of other codas, particularly given the highly limited morphological effects observed (see section 3.6.1.2).

3.6.1.2 Laxing Disharmony

Turning to Laxing Disharmony, we expected to find that intial-syllable laxing would be common in disyllabic words like /midi/ *midi* 'noon', which contain two identical high-vowel phonemes in open syllables. However, we find no significant effect supporting this process. Instead, it may be that the description of Laxing Disharmony stems from the observation that laxing may not necessarily be motivated by syllable shape and by the application of Laxing Harmony; tokens often surface as lax without an apparent trigger for laxing. However, future work should examine (a) whether a subset of speakers do favour Laxing Disharmony in speech just as they seem to in perception (Poliquin, 2006), and (b) whether other explanations exist for initial-syllable laxing in disyllabic words. For instance, Poliquin found that lexical frequency was not a significant

predictor when testing a small number of disyllabic words, but it may be that frequency effects influence the effect with a larger sample. From the current analysis, however, we cannot confirm the existence of Laxing Disharmony in the community-level phonological grammar.

3.6.1.3 Laxing Harmony

Regarding Laxing Harmony, the community-level grammar suggested by our results is complex. On one hand, we observe that high vowels are often significantly affected by the quality of adjacent vowels. As a general rule, speakers are more likely to produce a lax vowel when another high lax vowel is in an adjacent syllable. However, the effects are revealed to be far more complex when we examine the data in greater detail. The immediately preceding vowel typically has a significant effect such that (a) laxing is disfavoured when the vowel is not a high lax vowel, and (b) the effects appear to be associated with vowel height, other than for low vowels. We focus on the pattern in (a) here; section 3.6.2 will return to the gradient effect suggested by (b). Overall, our results suggest that base-final syllables have limited effects in conditioning harmony, contra expectations based on Fast (2008), though there is a consistent trend suggesting this base effect may simply be smaller compared to other motivations for harmony.

The overall patterns indicate the community-level phonological grammar a learner is expected to acquire. The final syllable has the strongest and most consistent effect overall, for all potential target syllables. This suggests that a learner will acquire Local Iterative Harmony in which neutral (i.e. non-high) vowels are transparent at the community level based on the final syllable's vowel predicting tenseness irrespective of intervening material. Furthermore, the robust effect of the quality of the vowel in the next syllable of the word confirms that Harmony is likely to pass through syllables iteratively. When we combine these results with the observation that the quality of the final-syllable vowel has a larger effect on the initial syllable of trisyllabic words than the quality of the vowel in the following syllable, we find evidence for both right-to-left Local Iterative Harmony and right-to-left Non-Local Harmony, which Poliquin (2006) inferred from the acceptability judgments of native speakers.

3.6.1.4 Lengthening Consonants

Recall that Walker (1984) observed that Laurentian French speakers lengthen vowels in non-final syllables when the intervocalic consonant (or consonant sequence) corresponded to a Lengthening Consonant. Based on this, we expected that Lengthening Consonants (aside from /r/) would trigger Pre-Fricative Tensing in non-final syllables. For example, we expected that the high vowels in the initial syllables of /vizit/ *visite* 'visit' and /viz+e/ *viser* 'to target' were likely to be longer and tense. This was only tested in disyllabic words, where we anticipated finding the strongest effects with the fewest potential confounds from other factors. The model for initial syllables of disyllabic words confirms that Pre-Fricative Tensing occurs in non-final syllables, but that morphological structure does not significantly modulate this effect. For example, *viser* is not more likely to undergo Pre-Fricative Tensing than *visite* even though the base of *viser* (/viz/) can be word-final and therefore undergo obligatory Pre-Fricative Tensing in morphologically related words (/viz/ *vise* 'targets' realised as [viːz]).

3.6.2 Phonetic and phonological effects

As previously suggested, it appears that laxing may partly result from coarticulatory (phonetic) pressures, rather than always reflecting a true (category-changing) phonological alternation. This was the focus of our second research question, which read as follows:

> Is laxing truly governed by the operation of phonological processes in non-final syllables or does it instead result from phonetic pressures? If laxing does result from phonetic pressures, is it phonetic reduction, coarticulation, or both?

The objective of this section is to discuss the evidence in favour of both phonological and phonetic accounts of laxing. To do so, we focus on two sources of evidence: the gradient effect of the height of vowels in adjacent syllables (section 3.6.2.1), and speech rate effects (section 3.6.2.2). We conclude that phonetic reduction does occur, but that phonological processes are primarily responsible for determining tenseness at the community level.

3.6.2.1 Gradience and Phonological Explanations of Laxing

In our analysis, we did not solely consider whether potential conditioning vowels for harmony were high lax vowels, even though descriptions of Laxing Harmony suggest that only high lax vowels can serve as triggers; we also examined whether the height of the vowel in adjacent syllables conditions laxing. If laxing were the result of coarticulation, then we would expect to find that lower vowels in adjacent syllables are associated with a higher rate of laxing in high vowels (i.e. gradient laxing rates based on height). Under this hypothesis, for example, the laxing rate of /u/ in the initial syllable would increase from /vuly/ *voulu* 'wanted.PST.PTCP' (high vowel) to /vule/ *voulez* 'want.IND.PRS.2PL' to /vulɛ/ *voulais* 'want.IND.IPFV.PST.3SG' to /vulɑ̃/ *voulant* 'want.PRS.PTCP'. We find that vowel height generally influences the rate of laxing even when a potential conditioning vowel is not both high and lax (i.e. is not a valid trigger for harmony). We find two main results in this regard. First, low vowels often pattern differently from other vowels by being associated with higher laxing rates than higher (non-high) vowels (e.g. mid low vowels compared to low vowels in disyllabic words). This effect can be analysed as phonological if mid high and mid low vowels, which are traditionally analysed as differing in height (see e.g. Lamontagne, 2014), instead differ in laxing (tense mid and lax mid, respectively). On the latter view, intervening low vowels would be underspecified for tense/lax and would therefore have neither height nor laxing features that could be visible to parasitic Laxing Harmony, while intervening mid vowels would bear values for tense/lax that could influence phonological computation and result in variably blocking Laxing Harmony. Second, laxing rates are not strongly correlated with vowel height, though there is a trend for increasingly lower neighbouring vowels to be associated with increasingly higher high-vowel laxing rates overall. This suggests a gradient coarticulatory effect because lowering the tongue body during a tense high vowel approximates the articulation of a lax high vowel (Dalton, 2011; Mielke and Dalton, 2013). The following section expands on this by considering the interaction between gradient coarticulation on one hand and speech rate on the other, based on our expectation that coarticulation would be magnified at faster speech rates due to articulatory ease.

3.6.2.2 Speech Rate and Phonological Explanations of Laxing

In our discussion of potential phonetic pressures that might result in lax vowels, we proposed two types of effects: phonetic reduction (undershoot) at faster speech rates and greater

coarticulation at faster speech rates. If phonetic reduction motivates laxing, then the /i/ in /arive/ *arrivé* 'arrived' and /arivɑ̃/ *arrivant* 'arriving' will both surface as lax more often in fast speech than in slow speech, even though there is no phonological trigger for laxing in either case. If coarticulation motivates laxing, then the /i/ in *arrivant* will lax a bit more often than the one in *arrivé* because the former contains a lower vowel in the following syllable – and this difference in laxing rate will be larger at faster speech rates because the effects of coarticulation are magnified. Overall, we find little in terms of speech rate effects; speech rate has a significant main effect in only one model (disyllabic words) and is only consistently present in significant or marginal interactions in two cases (disyllabic words, the penult of trisyllabic words). The expectation that faster speech on its own (i.e. without factoring in phonological context) may induce higher rates of laxing as a result of undershoot because the tongue does not reach the target position to articulate a high tense vowel cannot be confirmed.

The only case where speech rate is a main effect is in disyllabic words, but the interaction with vowel quality in the following syllable largely negates this effect. As a result, for these disyllabic words, we instead find a relatively categorical effect whereby producing a lax high vowel in the initial syllable is less likely in slower speech without a high lax vowel present in the final syllable. This suggests that phonological processes rather than phonetic reduction or coarticulation are at play because coarticulation with non-high vowels would likely result in more, rather than fewer, laxed realisations (lax high vowels are intermediate between tense high vowels and non-high vowels). Given that disyllabic words typically have higher lexical frequency than longer words and that speech rate correlates with lexical frequency in French (confirmed using the Lexique database; New at al. 2001, 2004), it may be that the effect found for disyllabic words is, in fact, a frequency effect only present in the most frequent words. This may, in fact, explain Laxing Disharmony observations in the literature: Laxing Disharmony may be favoured in highly frequent words like /midi/ *midi* 'noon' commonly being realised with a lax initial vowel despite the lack of apparent triggers for Laxing Harmony or Closed-Syllable Laxing. Frequency may also explain the variable results for base-triggered Laxing Harmony. A frequent morphologically complex form may be stored directly in the lexicon rather than being derived during production planning (i.e. accessing lexical representations, applying phonological alternations, and preparing motoric implementation), and consequently, morphological structure

may play a smaller role in laxing for those words (e.g. Henderson, 1985). We leave the question of frequency effects for future work.

In the case of high vowels in the penult of trisyllabic words, we have evidence for the application of phonological processes rather than phonetic pressures motivating the realisation of lax variants. Recall that these syllables are sensitive to the preceding (i.e. initial) syllable's tenseness rather than the final syllable's tenseness, which is unusual. For example, the /y/ in /buskylad/ *bousculade* 'hustle' or in /figyrin/ *figurine* 'figurine' is more likely to undergo laxing if the initial vowel surfaces as lax, with the final-syllable vowel having no statistically significant effect. Unexpectedly, laxing in the medial syllable after a high lax vowel is more common when speaking at *slower* speech rates. This suggests that speakers must plan laxing in the initial syllable (which appears to be Non-Local Harmony triggered by the final syllable) and must then plan laxing in the penult triggered by that initial syllable (Iterative Harmony applying left-to-right) in order to lax the medial vowel through harmony. At faster speech rates, however, the time available to plan this complex interaction is limited, and therefore speech rate effects may result from the difficulty of planning all phonological processes (e.g. Wagner, 2012; see also chapter 4). Since speech rate is significant in so few models, we conclude that the origin of lax high vowels in non-final syllables in Laurentian French cannot be attributed to phonetic pressures (reduction or non-phonologised coarticulation), but leave the examination of production planning effects for future work.

3.6.2.3 Summary of the Phonological Status of Laxing Processes

Based on our evidence from gradience and from speech rate effects, we conclude that lax high vowels are indeed largely the result of phonological processes, rather than phonetic pressures for reduction or coarticulation. Speech rate effects may in fact directly support the analysis that multiple (phonological) harmony processes apply because the interaction between vowel quality and speech rate in trisyllabic words is consistent with speakers needing to plan harmony as part of production planning. In summation, we have shown that phonological processes are predominantly responsible for determining the tense/lax quality of high vowels – with variation resulting both from phonetic pressures and from production planning limitations.

3.7 Conclusion

The current study provides the first large-scale quantitative analysis of high-vowel tenseness in non-final syllables in Laurentian French. It exploits the ability to train a forced aligner on obligatory laxing and tensing processes in final syllables to classify high vowels in non-final syllables. As such, it provides the first corpus of production data to test the complex interaction of processes affecting the tense/lax quality in non-final syllables. The first research question sought to determine which phonological processes a learner of Laurentian French would be expected to acquire based on their input, and therefore the phonological grammar of laxing that a learner would be expected to generate. The results confirm the presence of right-to-left Laxing Harmony of two types: Non-Local Harmony, and Local Iterative Harmony with transparent neutral vowels. We find additional evidence suggesting left-to-right Local Harmony, which should be further tested in future work given that it appears to be limited to trisyllabic words. However, we cannot confirm Closed-Syllable Laxing or Laxing Disharmony. The results further support optional Pre-Fricative Tensing in non-final syllables, even in words where obligatory Pre-Fricative Tensing does not occur in a final syllable at any stage of the derivation. They are additionally consistent with Laxing Harmony being opaque in the surface form. We argue that, in spite of massive complexity and variability in individual realisations, learners would be capable of generating a phonological grammar that includes specific Laxing Harmony processes (contra Poliquin, 2006).

Our second research question aimed to determine whether high-vowel tenseness is predominantly governed by category-changing phonological processes or whether phonetic pressures accounted for a considerable proportion of lax high vowels in Laurentian French. From the overall lack of speech rate effects and the limited gradient effects conditioned by the height of vowels in adjacent syllables, we conclude that laxing reliably involves category-changing phonological processes and is not only the result of phonetic pressures.

We leave for future work two main questions. First, we do not address individual differences in the current study (aside from the inclusion of random intercepts to control for distinct baseline rates), instead focusing on community-level patterns that learners are expected to acquire. Data exploration suggests that, for disyllabic words, speakers converge towards a shared laxing rate of about 12% in the initial syllable. Speakers providing more tokens have laxing rates for the initial

syllable of disyllabic words that are closest to this laxing rate. While this suggests that speakers may not differ as much as previously proposed by Poliquin (2006), who suggested that non-laxing speakers exist, it does not address whether different speakers employ different harmony processes that reach this same rate of laxing overall. An examination of the random effects in our models suggests that Poliquin overestimated the differences between speakers; variance associated with words was consistently greater than variance associated with speakers. As such, future work should investigate whether individual speakers differ in their grammars, like Poliquin suggests his twelve speakers do from the acceptability judgments obtained.

Second, future work should further investigate Closed-Syllable Laxing and High-Vowel Deletion. In the current study, following Poliquin (2006), we syllabified based on surface forms and thereby treated the onset of a syllable with a deleted nucleus (e.g. the medial underlying syllable in /difisil/ *difficile* 'difficult' when realised as [dʑɪfsɪl]) as closed. However, future studies should test whether these syllables instead pattern like open syllables, or like neither open nor closed syllables. This analysis, we believe, is best conducted at the level of individuals because speakers may not all share the same analysis for these derived cases.

We conclude with a methodological note, which we believe is the most wide-reaching contribution of the current study. In some cases, like for high-vowel tenseness in Laurentian French, acoustic similarity on even a small number of acoustic dimensions (e.g. the first two formants) prohibits large-scale quantitative analyses of production. However, exploiting ever-improving tools such as forced aligners can render previously inconceivable analyses or untestable problems considerably more plausible. We hope that the methods outlined in the current study will lead to more studies that tackle problems requiring large datasets due to complexity or variability. In the current study, for instance, those methods allowed us both to show that speakers could generate a grammar of laxing (contra the claim by Poliquin, 2006) and what that grammar is expected to be.

References

ANQ. (2011). *Assemblée nationale du Québec*. http://www.assnat.qc.ca/fr/index.html. Website of Assemblée nationale du Québec.

Arnaud, V., C. Sigouin, and J.-P. Roy. (2011). Acoustic description of Québec French high vowels: First results. In *Proceedings of the International Conference on Phonetic Science* 17: 17-21.

Baronian, L. (2001). *Directional vowel harmony: The case of Quebec French*. Paper presented at TREND 11 (Stanford University, April 2001).

Bates, D., M. Maechler, B. Bolker, S. Walker. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67(1): 1-48.

Bosworth, Y. (2011). *Weight and feet in Québécois*. PhD thesis, University of Texas at Austin.

Côté, M.-H. (2008). *La syllabation et le relâchement des voyelles hautes en syllabe non-finale en français québécois*. Paper presented at Journées PFC, Tulane University.

Cote, M.-H. (2010). La longueur vocalique devant consonne allongeante en contexte final et derive en francais laurentien. In: C. LeBlanc, F. Martineau and Y. Frenette, (eds.), *Vues sur les français d'ici*, Quebec City: Presses de l'Universite Laval, 49–75.

Côté, M.-H. (2012). Laurentian French (Québec): Extra vowels, missing schwas and surprising liaison consonants. In: R. Gess, C. Lyche, and T. Meisenburg, (eds.), *Phonological Variation in French: Illustrations from Three Continents*, Amsterdam: John Benjamins, 235-274.

Côté, M.-H., and H. Saint-Amant Lamy. (2012). D'un [ɾ] à l'aut[ʁ]e: Contribution à la chute du R apical au Québec. In: F. Neveu, V. Muni Toke, P. Blumenthal, T. Klinger, P. Ligas, S. Prévost, and S. Teston Bonnard, (eds.), *3ème Congrès Mondial de Linguistique Française*, Paris: Institut de Linguistique Française, 1441-1453.

Cuerrier, A., and C. Reisss. (2018). Geminates and vowel laxing in Quebec French. *Proceedings of LSRL 47*.

Dalton, W. (2011). An ultrasound imaging study of the tense-lax distinction in Canadian French vowels. *Canadian Acoustics* 39(3).

Déchaine, R.-M. (1990). *Stress and weight gain*. Ms., University of Massachusetts, Amherst.

Dumas, D. (1983). *Nos façons de parler*. Québec: Presses de l'Université de Québec.

Dumas, D. (1987). *Nos façons de parler: Les prononciations en français québécois*. Québec: Presses de l'Université de Québec.

Durand, J., B. Laks, and C. Lyche. (2002). La phonologie du français contemporain: usages, variétés et structure. In: C. Pusch, and W. Raible, (eds.), *Romanistische Korpuslinguistik - Korpora und Gesprochene Sprache*, Tübingen: Gunter Narr Verlag, 93-106.

Fast, A. (2008). Optimality and opacity in canadian french vowel harmony: A variationist account. *Revue des étudiants en linguistique du Québec* 3(1).

Flemming, E. S. (2005). *A Phonetically-Based Model of Phonological Vowel Reduction*. Ms., MIT.

Garcia, Guilherme D., Heather Goad and Natália Brambatti Guzzo. (2017). Footing is not always about stress: Formalizing variable high vowel deletion in Québec French. In: K. Jesney, C. O'Hara, C. Smith, and R. Walker, (eds.), *Proceedings of the Annual Meeting on Phonology 2016*.

Grammont, M. (1914). *Traité de prononciation française*. Paris: Delagrave.

Hannahs, S. (1989). High vowel variation in Québec French. In *Annual Meeting of the Linguistic Society of America*, Washington, DC.

Henderson, L. (1985). Towards a psychology of morphemes. In: A. Ellis, (ed.), *Progress in the Psychology of Language*, vol. 1, London: LEA, 15-72.

Lamontagne, J. (2014). Examining opposing processes in Laurentian French vowels. *Proceedings of the 2014 annual conference of the Canadian Linguistic Association*. http://cla-acl.ca/?p=1188

Légaré, L. (1978). Relâchement des voyelles hautes et reformulation. *Cahier de linguistique* 7: 31-42.

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

Mielke, J., and W. Dalton. (2013). *Many-to-one articulatory-acoustic mappings in Canadian French vowels*. Ultrafest VI, University of Edinburgh.

Milne, P. (2014). *The variable pronunciations of word-final consonant clusters in a force aligned corpus of spoken French*. PhD thesis, University of Ottawa.

New, B., C. Pallier, L. Ferrand, and R. Matos. (2001). Une base de données lexicales du français contemporain sur internet: Lexique. *L'Année Psychologique* 101: 447–462. http://www.lexique.org.

New, B., C. Pallier, M. Brysbaert, and L. Ferrand. (2004). Lexique 2: A new french lexical database. *Behavior Research Methods, Instruments, and Computers* 36(3): 515–524.

Noske, R.G. (1982). Syllabification and syllable changing rules in French. In: H.G. van der Hulst, and N.S.H. Smith, (eds.), *The structure of phonologicalrepresentations*, vol. 2. Dordrecht: Foris, 257-310.

Poliquin, G. (2006). *Canadian french vowel harmony*. PhD thesis, Harvard University.

R Development Core Team. (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, http://www.R-project.org.

Sigouin, C. (2013). *Caractéristiques acoustiques des voyelles fermées tendues, relâchées et allongées en français québécois*. MA memoire, Université de Québec à Chicoutimi.

Sigouin, C. and V. Arnaud. (2015). Québec French close vowels in lengthening contexts: Tense, lax or diphthongised? an acoustic study. In: *Proceedings of the International Congress on Phonetic Science 2015*, Glasgow.

Spilková, H. (2014). *Phonetic reduction in spontaneous speech:an investigation of native and non-native production.*. PhD Thesis, Norwegian University of Science and Technology – Trondheim.

Wagner, M. (2012). Locality in phonology and production planning. *McGill Working Papers in Linguistics* 22(1): 1-18.

Walker, D. (1984). *The pronunciation of Canadian French*. University of Ottawa Press.

Preface to Chapter 4

Forced alignment is not only a tool for data preparation by segmenting audio files, but also a tool that can perform data classification to better understand the structure of the data. Large-scale corpora render forced alignment possible; a substantial amount of data is required to train accurate phone models. Sizeable corpora are also required for detailed investigation of phenomena in which multiple factors determine phonetic and phonological realisations or in which factors have subtle effects.

Chapter 3 examined the tense/lax quality in high vowels using corpus data from Laurentian French. The use of automated classification through forced alignment made it feasible to test the many – typically optional – phonological processes that have been proposed to determine the tense/lax quality in non-final syllables. The results show mixed support for processes proposed in the literature: some, like transparent Laxing Harmony and Retensing from Resyllabification, are supported, while others, like Laxing Disharmony, are not. In spite of the multiple processes involved, Chapter 3 demonstrates that there is sufficient input for learners of Laurentian French to acquire the complex phonology that determines the tense/lax quality in high vowels.

Together, Chapters 2 and 3 demonstrated that variation is sensitive to a range of factors such as the phonological environment, morphological structure, and prosodic position. Speakers must therefore cognitively process a considerable amount of information to apply phonological processes – in addition to retrieving words and phonologically encoding segments. The ease of planning words and their contexts could therefore have a considerable effect on process application rates, particularly for processes that apply across words and therefore require retrieving more than one word. Chapter 4 tackles the problem of cognitive processing by testing the hypothesis that production planning limitations result in variability. Based on the literature, it is expected that production planning correlates with the direction of application of phonological processes. Cross-word hiatus resolution in Madrid Spanish serves as an ideal test case because V1 or V2 can be targeted for deletion in /V1#V2/ sequences. Chapter 4 therefore probes

directional asymmetries in hiatus resolution in Madrid Spanish to determine how cognitive processing influences variability.

Chapter 4

Cross-word hiatus in Madrid Spanish

4.1 Introduction

Predictability effects in speech have been well-documented in phonetic and phonological research for almost a century, with more predictable words and word sequences consistently shown to undergo greater levels of reduction (e.g. Jespersen, 1922; Zipf, 1929; Lindblom, 1983, 1990; Moon and Lindblom, 1994; Sotillo, 1997; Bard and Aylett, 1999; Gregory et al., 1999; Aylett, 2000; Bard et al., 2000; Bush, 2001; Jurafsky et al., 2001; Bell et al., 2003). The main explanation for this finding lies in cognitive processing: lexical retrieval is faster for words that are frequent or predictable (e.g. Bell et al. 2009). The Locality of Production Planning Hypothesis (henceforth PPH; Wagner, 2011, 2012; Tanner et al., 2015; Kilbourn-Ceron, Wagner and Clayards, 2016; MacKenzie, 2016; Kilbourn-Ceron, 2017; Tamminga, 2018; Bailey, 2019) draws on findings from this literature on phonetic reduction, but expands upon the proposal that (variable) reduction is a consequence of lexical retrieval. The PPH states that cross-word (sandhi) processes are inherently variable because the narrow window available for phonological encoding limits the ability to plan requisite contexts and therefore the limits the application of cross-word processes. In this respect, the PPH departs from previous frequency-based accounts in focusing on how cognitive processing affects phonological operations that may apply across word boundaries, rather than on how ease of lexical retrieval affects the realisation of individual words. Crucially, this means that a word's frequency and conditional probability can be associated with the rate of application of a process targeting *another* word. For instance, Kilbourn-Ceron, Wagner and Clayards (2016) demonstrate that the rate of flapping of coronal stops across words in English (e.g. the /t/ in *but everyone* being realised as [ɾ] because *everyone* provides the following vowel that triggers flapping) is significantly correlated with the lexical frequency of the upcoming word (*everyone*).

In effect, the PPH shifts the focus from the word or segment that is targeted by a process to the context required for that process because the context (trigger) must be planned early enough for the target to undergo the process. One implication of this is that there should be a directional asymmetry, as discussed by Kilbourn-Ceron (2017). This, however, has yet to be tested empirically. Hiatus resolution offers a particularly effective test case for the PPH in languages where either V1 or V2 in a /V1#V2/ sequence can be the target of hiatus resolution processes like deletion. In right-to-left (target before trigger) hiatus resolution, we predict that deletion of V1 will be sensitive to production planning effects because the trigger (V2) must be planned early enough for V1 to be deleted. In left-to-right (trigger before target) hiatus resolution, however, deletion of V2 is unlikely to be modulated by production planning because the trigger (V1) is expected to already have been planned due to its earlier position. If both V1 and V2 can be targeted by deletion, then either vowel can act as trigger or target and therefore left-to-right and right-to-left applications can be directly compared. As such, hiatus where a vowel in either position can delete allows us to test the prediction that processes relying on upcoming material (i.e. right-to-left processes) will be more sensitive to production planning factors than processes that do not (e.g. left-to-right processes). This is because the window for phonological encoding is expected to include the process trigger in the left-to-right case, but this is less certain in the right-to-left case.

Spanish has a hiatus resolution process like the one exemplified here, meaning it offers an optimal test case for the directional asymmetry inherent to the PPH. Underlying vowels in /V#V/ sequences can optionally be deleted or reduced to a glide in Spanish (Aguilar, 1999; Hualde and Chitoran, 2003; Baković, 2006; Vuskovich, 2006; Souza, 2010). The literature has predominantly focused on word-internal sequences, consistently finding a preference to preserve more sonorous vowels over less sonorous ones (ex. /a/ over /e/).[39] In the current study, we will focus on Madrid Spanish, where we show deletion to be a common method of resolving instances of cross-word hiatus.

[39] The current study does not probe the motivations for /e/ being the primary target of deletion. We do, though, point out that /e/ is both the most frequent phoneme in spoken Spanish (Moren Sandoval et al., 2008) and the epenthetic vowel in Spanish (Eddington, 2001), and therefore motivations like its high recoverability may contribute to its relative ease of deletion.

We use the PraatAlign forced aligner (Lubbers and Torreira, 2016) to categorise underlying vowel-vowel sequences across words in the Nijmegen Corpus of Casual Spanish (Torreira and Ernestus, 2010), allowing us to determine whether the vowel sequence was realised as V1, V2, or V1V2 in spontaneous speech. We then analyse how factors related to production planning (lexical frequency, conditional probability[40]) predict the resulting classifications (V1, V2, V1V2) using mixed-effects multinomial regression. Our analysis confirms that /a/ deletion is relatively rare and further demonstrates that the asymmetry predicted by the PPH is borne out in the data: /e/ deletion is more sensitive to production planning factors in /e#a/ (target first) than in /a#e/ (trigger first), consistent with the prediction that ease of planning phonological contexts mediates process application rates.

4.2 Methods

In the current study, we analyse 2,426 tokens of the vowel sequences /a#e/ and /e#a/. We use mixed-effects logistic regression in R (R Development Core Team, 2015) to predict whether each vowel in the sequence will be deleted or not (since it could be that neither is) in order to explore predictability effects affecting vowel deletion. In section 4.2.1, we discuss the corpus employed in this study and the alignment of the corpus recordings. We continue in section 4.2.2 by describing the data extraction, the statistical analysis, and our predictions.

4.2.1 Corpus and alignment

The data examined in this study come from the Nijmegen Corpus of Casual Speech (Torreira and Ernestus, 2010), a corpus composed of spontaneous conversations between speaker pairs and triads who knew each other before the corpus was collected. The use of casual spontaneous speech is essential for our purposes, as previous research has found that hiatus resolution in Spanish occurs significantly less frequently in read speech than in spontaneous speech (Aguilar, 1999). There are two main explanations for differences in application rates between read and spontaneous speech that are of interest for the current study. The first explanation is that hiatus

[40] These factors may also influence a listener's capacity to correctly parse the surface realisation despite the application of vowel deletion. The question of interpretability will be discussed in section 4.3.

resolution may be affected by speech formality or other sociolinguistic considerations, which often suppress variation in more formal contexts (Labov, 1972). Consistent with this trend for register differences to influence variation, Aguilar (1999) finds that deletion is more frequent in Spanish hiatus contexts in spontaneous speech compared to read speech. The second explanation of interest is that there are differences in lexical retrieval between read and spontaneous speech that result from retrieving a word presented to the speaker – sometimes without needing to process semantic content or syntactic structure – rather than accessing the word from the lexicon as part of normal production planning (Navarrete et al., 2015). It is therefore beneficial to examine spontaneous speech when testing the PPH, given that we want to minimise the risk of speakers inhibiting hiatus resolution and given that lexical retrieval plays an important role in determining the window of phonological encoding.

Recordings from 54 speakers were aligned using the PraatAlign forced aligner (Lubbers and Torreira, 2016) with Spanish phone models obtained from the Munich Automatic Segmentation System (MAUS) project (Schiel, 1999). Forced alignment requires a minimum duration for each segment. For PraatAlign, for example, the aligner uses tri-state models for each vowel, which equates to looking for three consecutive slices for each vowel. Since the recording is parsed by the aligner in 10-millisecond slices, the minimum duration for a single segment is 30 milliseconds (three 10-millisecond slices). This is problematic for our purposes because we are examining the deletion of a segment in a sequence of two underlying segments, which means that the aligner would be trying to match six target states and would therefore require a minimum duration of 60 milliseconds (two segments with three 10-millisecond slices each), when only one segment may actually be produced. This minimum duration would inevitably lead to alignment errors in exactly the context that interests us, as we examine contexts where one vowel – rather than two – may be pronounced. As a result, we only include /V#V/ sequences of at least 60 milliseconds in the current analysis. Based on typical durations for individual vowels (see Stevens, 2011), we are confident that 30 milliseconds per aligned segment is a reasonable threshold for vowel sequences in the current study.[41]

[41] Preliminary work on the 2,136 tokens with durations of less than 60 milliseconds suggests that they exhibit the same patterns as longer tokens (Lamontagne and Torreira, 2017). However, because the tokens must be analysed separately due to different possible surface realisation options (short tokens cannot be aligned as having two surface vowels), we focus on the tokens of at least 60 milliseconds in this analysis. The resulting sample size is the aforementioned 2,426 tokens, reduced from 5,060.

The main advantage of PraatAlign for our study is that it allows users to specify phonological rules a priori (e.g. context-sensitive vowel deletion). In our case, we instructed the aligner that, when looking for two consecutive vowels across a word boundary, only one vowel might be produced. In *la escuela* 'the school', for example, the aligner considers three possible pronunciations for the /a#e/ sequence: [ae], [a] and [e].[42] As such, the aligner had the possibility of aligning the underlying vowel sequence as having undergone deletion, but had no built-in bias regarding process application or preferred target vowel. While automated classification remains a relatively new method in phonetic and phonological research, it has proven effective for analysing sound patterns in spontaneous speech (e.g. Schuppler et al., 2011; Schuppler et al., 2014; see also Chapter 3).

4.2.2 Procedure and predictions

Once the recordings were aligned, 2,426 tokens of /e#a/ and /a#e/ of at least 60 milliseconds with no intervening pauses, as determined by the forced aligner, were measured and extracted by script using Praat (Boersma and Weenink, 2018). We focus on /a#e/ and /e#a/ for three main reasons. First, using non-identical vowels allows us to be certain of the vowel identity if only one vowel is preserved. Second, the literature (e.g. Garrido, 2008) is most clear about the hiatus resolution preferences for /e/ (which deletes relatively readily) and /a/ (which is less likely to delete). Crucially, this pairing includes vowels with different deletion preferences, which enables us to compare cases where the preferred target for deletion is clear. Third, neither vowel is known to undergo gliding in the dialect.[43] The resulting tokens are well-balanced, with 51.8% of tokens coming from /e#a/ sequences and the remaining 48.2% coming from /a#e/ sequences, and both vowels occurring in function and content words.

[42] Note that the aligner does not label slices directly, determining that two slices from /a/ are followed by four slices from /e/, for instance. Instead, the aligner evaluates probabilities of an alignment; if the aligner determines that the acoustic signal is an 87% match for /a/ and only a 30% match for /e/ and a 46% match for /ae/, it will align that portion of the acoustic signal as containing /a/.

[43] The information that non-high vowels do not typically undergo gliding has not yet been verified empirically to the authors' knowledge for Madrid Spanish speakers, but is obtained predominantly from the intuitions of the second author, who is a native speaker familiar with the dialect and who collected the corpus. Future work should investigate this possibility further, but cursory examinations of the corpus are consistent with this observation.

In this paper, we examine the automatic classification obtained from PraatAlign, coding for whether the first vowel, the second vowel, or neither vowel was deleted. We employ Bayesian mixed-effect multinomial regression using the MCMCglmm package (Hadfield, 2010) in R, with the word and speaker as random intercepts to control for individual differences as well as lexical effects (e.g. stress and word length). We performed 10,300 iterations with 1,000 samples per iteration and 300 iterations were processed before samples were stored. Two types of factors were examined in the analysis: (1) segment-level factors, (2) factors related to production planning. First, we examine segmental factors, which are expected to interact with production planning factors. We consider the effects of **underlying vowel sequence** (/a#e/ vs. /e#a/). Based on previous research, we predict that /a/ will be much less likely to delete than /e/ (Garrido, 2008) and therefore we anticipate that we are effectively comparing the deletion of /e/ in /a#e/ and /e#a/ sequences. As a result of a cross-linguistic preference to preserve word-initial segments (e.g. Barnes, 2006 and Keating, 2006 for reviews), we predict that deletion of V1 will be favoured over deletion of V2 in /V#V/ sequences. Taken together, this implies that V1 realisations will be preferred more for /a#e/ than for /e#a/, and that V2 realisations will be preferred more for /e#a/ than for /a#e/. It additionally implies that /e#a/ is more likely to be realised as [a] than /a#e/ is, since the preference to preserve /a/ and the preference to retain word-initial segments favour the same surface form only in /e#a/.

Second, we consider production planning factors. These are factors that seek to capture whether the speaker had planned the underlying vowel hiatus sequence and therefore deleted the first vowel in anticipation of a preferred vowel (in our case, /a/). Here we consider the **lexical frequencies** of both words (i.e., p(word 1) and p(word 2)) and the **conditional probabilities** of each word given the other (i.e., p(word 1|word 2) and p(word 2|word 1)), which allows us to effectively estimate how likely the word is in general and how likely it is in context, both of which are associated with faster lexical retrieval and phonological encoding during production (e.g. Segui et al., 1982; Wagner, 2012). The conditional probabilities were calculated based on the bigram frequencies obtained from a 20-million-word sub-corpus of the Corpus del Español Actual (Subirats and Ortega, 2012), in combination with the 300,000-word Nijmegen Corpus of Conversational Spanish. Both lexical frequency and conditional probability were log-transformed.

The **interaction** between production planning factors and the vowel sequence is also considered: if speakers delete /e/ more often than /a/, we expect that having planned that V2 is /a/ will increase the probability of deleting /e/ in V1. In the design of this study, this entails an asymmetry whereby V1 deletes more often than V2 in /e#a/ sequences, but V2 deletes more often than V1 in /a#e/ sequences. This difference should then be mediated by lexical frequency and conditional probability; the ability to plan upcoming material is more important when it facilitates planning a trigger for deletion (deleting /e/ in /e#a/) compared to when that trigger appears earlier (deleting /e/ in /a#e/). In short, for both words, higher lexical frequency and higher conditional probability are predicted to increase /e/ deletion rates in /e#a/ sequences more than in /a#e/ sequences because those factors predict the speaker's ability to plan the context for /e/ deletion and therefore apply /e/ deletion more often.

We additionally control for word class as a proxy for prosodic structure. Preliminary analysis suggests that the word class (content or function) of the first and second word is predictive of prosodic structure and therefore that including word classes as predictors provides a control for prosodic structure (Lamontagne and Torreira, 2017). For this study, function words have been defined as words belonging to a closed class (e.g. determiners and prepositions), while content words are those belonging to an open class (e.g. nouns and verbs). Control of prosodic structure is essential in the current study for two reasons. First, prosodic structure is independently associated with production planning, although not all prosodic effects can be attributed to production planning (Kilbourn-Ceron, Wagner and Clayards, 2016; Kilbourn-Ceron, 2017). Second, function words typically have higher frequency than context words, which results in function words typically having faster lexical retrieval than content words (e.g. Segalowitz and Lane, 2001). Consequently, conditional probabilities are correlated with syntactic structure, which in turn is associated with prosodic structure. It is relatively unlikely that there is a strong prosodic boundary between a function word (word 1) and a content word (word 2); for example determiners precede nouns in Spanish (e.g. *la escuela* 'the school') and the two form a single prosodic unit (phonological phrase). Testing lexical frequency and conditional probability without controlling for prosodic structure could therefore yield confounded results; both factors become predictors of word class and syntactic/prosodic structure, rather than of production planning more broadly. In short, we control for prosodic structure using word class, but we do

not focus on the effects of word class in the current study because our primary interest is the role of production planning (but see Chapter 5 for related discussion).

4.2.3 Summary

In short, we use automated classification of Spanish hiatus sequences across words to compare factors capturing a speaker's ease of production planning. We predict that the effect of production planning factors will crucially be asymmetrical: these factors will be most important when they reflect the ability to plan the phonological context for deletion (target-first; /e#a/ sequences), but will have little or no significant effect when the context for deletion is known from the previous word (trigger-first; /a#e/). In the next section, we address this interplay as evidenced by our data.

4.3 Results

We here examine the results from a mixed-effect multinomial regression of 2,426 underlying /e#a/ or /a#e/ sequences, predicting whether the surface form is V1V2 (no deletion), V1 (V2 deleted) or V2 (V1 deleted). The model output is provided in Table 18. Each predictor in the model is shown alongside the surface form (V1 or V2) that is favoured or disfavoured by the predictor. The model treats the surface form without deletion (i.e., V1V2) as its default value, which is why it is not included in the table like the two other surface forms are. With respect to the model selection parameter for the output below, the deviance information criterion (DIC) was equal to 3217.155.

We begin by describing the general patterns in Madrid Spanish based on vowel quality and position in the word (section 4.3.1). We then turn to the production planning effects of interest: frequency (section 4.3.2), followed by conditional probability (section 4.3.3). In all cases, our primary goal is to compare left-to-right (V2) and right-to-left (V1) deletion, where we predict production planning factors to most strongly mediate right-to-left deletion. This goal is the main focus of section 4.3.4, where we summarise our results.

Throughout this section, we include visualisations to facilitate interpretation of our results. The figures are created directly from the data analysed in our models, therefore including only tokens with durations of at least 60 milliseconds. Because our model considers the effects of our

controls (word classes) in addition to our factors of interest (underlying vowel sequence, lexical frequency, and conditional probability), the figures often understate effect sizes and confidence intervals relative to our model. As a result, we focus predominantly on our model output to guide our interpretation.

Table 18: Model output for realisation patterns in /a#e/ and /e#a/ sequences.

	Posterior mean	Lower 95% CI	Upper 95% CI	Effective sample	p (MCMC)	
[V1]	2.464424	1.274597	3.836742	54.53	<0.0001	***
[V2]	1.099917	-0.18811	2.408472	60.52	0.0864	.
[V1]:/e#a/	-1.74397	-3.79404	0.266196	31.44	0.0866	.
[V2]:/e#a/	3.462651	1.551061	5.513728	39.49	<0.0001	***
[V1]:w1 class	0.378442	-0.11106	0.879762	60.3	0.1558	
[V2]:w1 class	-0.1883	-0.70567	0.316694	79.84	0.4854	
[V1]:w2 class	0.491591	-0.02063	0.975219	43.78	0.0634	.
[V2]:w2 class	0.901456	0.294688	1.518202	41.54	0.0048	**
[V1]:p(w2|w1)	-0.00081	-0.05162	0.050326	57.73	0.9246	
[V2]: p(w2|w1)	-0.01624	-0.08163	0.059073	62.41	0.5702	
[V1]:p(w1|w2)	0.15031	0.037449	0.263954	46.59	0.0188	*
[V2]:p(w1|w2)	0.232209	0.083941	0.381906	25.81	0.001	***
[V1]:p(w1)	0.109648	0.029489	0.190424	75.32	0.0058	**
[V2]:p(w1)	0.145809	0.045189	0.2497	65.48	0.002	**
[V1]:p(w2)	-0.07351	-0.19472	0.040427	51.29	0.201	
[V2]:p(w2)	-0.20645	-0.3473	-0.08666	52.45	<0.0001	***
[V1]:w1 class:w2 class	-0.35482	-1.66483	1.217334	52.66	0.5842	
[V2]:w1 class:w2 class	0.298627	-1.30122	2.019416	61.08	0.7238	
[V1]:/e#a/:p(w2|w1)	-0.0247	-0.09464	0.042115	26.64	0.4528	
[V2]:/e#a/:p(w2|w1)	0.022809	-0.03684	0.08155	39.53	0.44	
[V1]:/e#a/:p(w1|w2)	-0.05591	-0.27084	0.143633	33.44	0.6146	
[V2]:/e#a/:p(w1|w2)	0.014433	-0.19243	0.199357	46.3	0.8484	
[V1]:/e#a/:p(w1)	-0.14442	-0.37498	0.040319	22.25	0.119	
[V2]:/e#a/:p(w1)	-0.1823	-0.37947	-0.00043	35.95	0.0298	*
[V1]:/e#a/:p(w2)	-0.06274	-0.23542	0.121275	26.79	0.469	
[V2]:/e#a/:p(w1)	0.112431	-0.04495	0.290261	37.16	0.186	

4.3.1 Vowel deletion in Madrid Spanish

We begin with the overall patterns for vowel deletion in Madrid Spanish. Regardless of the vowel targeted by deletion, we observe that deletion of at least one vowel is the norm. We first consider vowel quality and the linear order of vowels. We predicted that /e/ would be more likely

to delete than /a/, meaning that V1 would delete more often in /e#a/ sequences and that V2 would delete more often in /a#e/ sequences. As illustrated in Figure 12, which plots vowel realisations, this is indeed what we find: /a/ deletes in 11% of sequences, while /e/ deletes in 77% of sequences. For /a#e/ (the baseline value in our model), deletion of V2 is favoured ($p<0.0001$) more strongly than deletion of V1 ($p=0.0864$). For /e#a/, however, the pattern is reversed; deletion of V1 is favoured ($p<0.0001$). There is additionally a trend to delete V2 less often than for /a#e/ ($p=0.866$).

Figure 12: Surface forms by underlying vowel sequence.

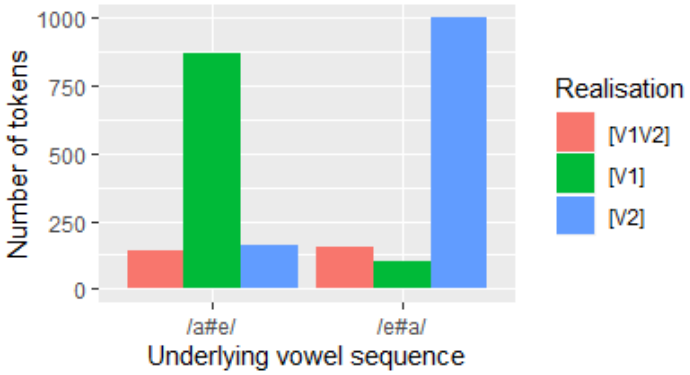

4.3.2 Lexical frequency and vowel deletion

We turn now to the effect of lexical frequency, as illustrated in Figure 13 (frequency of word 1) and Figure 14 (frequency of word 2). We predicted that vowels in frequent words would delete more often, consistent with the existing literature finding that reduction is sensitive to frequency. Our results are mostly consistent with this prediction. While deletion of V1 is more likely to occur when word 1 is frequent ($p=0.002$), deletion of V2 is not significantly associated with the frequency of word 2 ($p=0.201$). Following the PPH, we also predicted that a vowel's deletion rate would be sensitive to the other word's lexical frequency because the speaker must plan *both* words in order for deletion to apply. Consistent with this prediction, we find that V1 deletes significantly more often when word 2 is more frequent ($p<0.0001$) and that V2 deletes significantly more often when word 1 is more frequent ($p=0.0058$).

This support for the PPH increases when we consider the underlying vowel sequence. We predicted that right-to-left (/e#a/) deletion would be more sensitive to production planning factors like lexical frequency because the window for phonological encoding will only variably

include both V1 and V2. If only V1 is planned, then vowel deletion will not occur because the trigger for deletion (V2) is not yet planned. When the trigger comes first, however, it is expected that there will be little or no variability in phonological encoding, in which case production planning factors are not predicted to play a significant role. This is indeed what we find, as shown by the deletion rate of V1 in /e#a/ sequences being significantly more affected by the lexical frequency of word 1 (p=0.0298).[44] In short, higher lexical frequency is associated with greater ease to delete the first vowel, which is consistent with predictions based on the Production Planning Hypothesis.[45]

Figure 13: Deletion rates for /a#e/ and /e#a/ sequences by frequency of word 1.

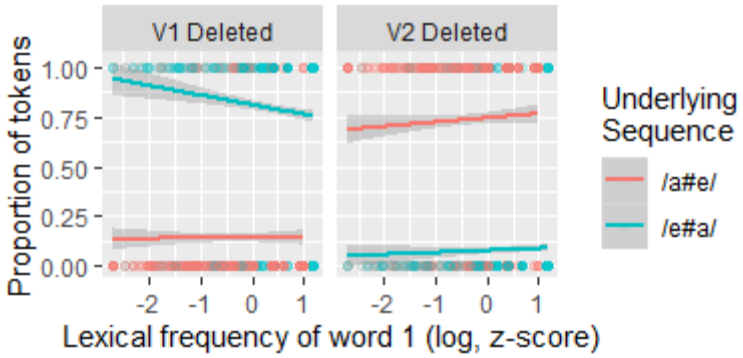

Figure 14: Deletion rates for /a#e/ and /e#a/ sequences by frequency of word 2.

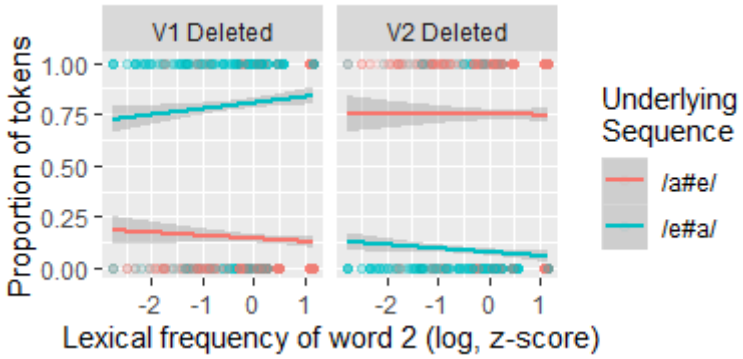

[44] Preliminary data exploration suggests that V1 deletion in /a#e/ sequences is predominantly favoured in cases when /a/ is both recoverable from context and results from earlier morphological planning, as in the case of the feminine suffix /-a/ that is marked elsewhere in the phrase through agreement. See Lamontagne and Torreira (2017) for an account of these cases.

[45] In Figure 13, we observe what appears to be a marginal trend for decreased /e/ deletion rates in /e#a/ when the first word is especially frequent. In the data under examination, this is largely restricted to certain lexical items that have especially high frequency and that are monosyllabic. This effect may result from the first word being planned too early for word 2 to also be planned, in which case the vowel is not deleted in order to provide more planning time. Alternatively, as proposed by Tanner et al. (2017) and consistent with the PPH, word 1 being planned so much earlier than word 2 would mean that the context for deletion is not determined early enough for deletion to apply.

4.3.3 Conditional probability and vowel deletion

Finally, we examine the effect of conditional probability, as illustrated in Figure 15 (probability of word 1 given word 2) and Figure 16 (probability of word 2 given word 1). Like for lexical frequency, we predicted that higher conditional probability would be associated with increased deletion rates due to the increased ease of planning upcoming material. Whereas lexical frequency showed relatively clear asymmetries between vowel quality and direction of deletion, this is not as clearly the case for conditional probability.[46] We only have high confidence that higher conditional probability of word 1 given word 2 is associated with higher deletion rates both of V1 (p=0.001) and of V2 (p=0.0188). A comparison of posterior means and of confidence intervals suggests that this effect is likely to be greatest for the deletion of V1, consistent with our predictions. Examination of Figure 4 and comparison of confidence intervals in our model suggest that the effect of conditional probability is probably larger for V1 deletion in /e#a/ than for V2 deletion in /a#e/, but with an especially high degree of variability. This comparison is the one of most importance for the current study; this is where the directional asymmetry implicit to the PPH is most clear. As such, while the evidence is limited – in part, we suspect, due to our sample size – we cautiously interpret this result to be consistent with the Production Planning Hypothesis. Future work should aim to test conditional probability effects further, in particular to determine whether lexical frequency effects are more robust overall.

Figure 15: Deletion rates for /a#e/ and /e#a/ sequences by the probability of word 1 given word 2.

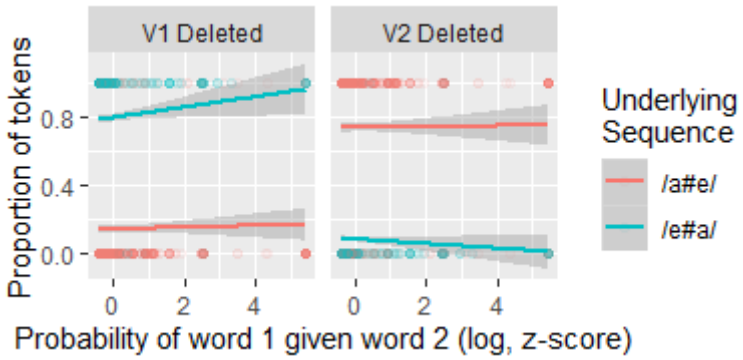

[46] As noted in section 2.2, this reduced effect may be due in part to the inclusion of word class in our model. However, we opted to include this factor to have greater confidence in the interpretation of conditional probability as resulting from production planning effects rather than from the prosodic structure with which they are associated.

Figure 16: Deletion rates for /a#e/ and /e#a/ sequences by the probability of word 2 given word 1.

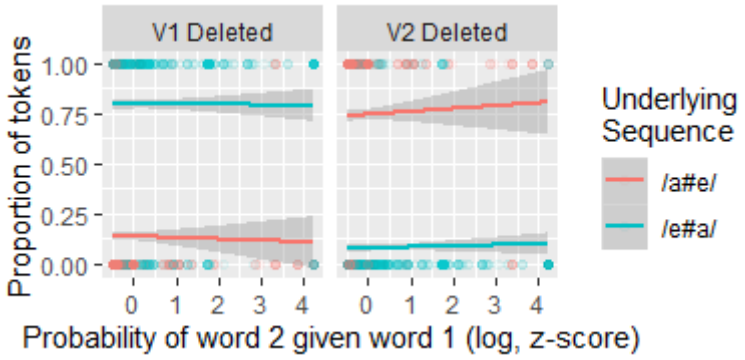

4.3.4 Summary

Overall, we find that hiatus resolution through deletion is exceptionally common in Madrid Spanish, occurring in nearly 88% of tokens. Deletion can target V1 or V2 in both /e#a/ and /a#e/ sequences. However, V1 deletion dominates in /e#a/ sequences, while V2 deletion dominates in /a#e/ sequences, indicating an overall preference to favour preserving /a/ over /e/. Production planning factors further influence this process by mediating deletion rates; the more easily a speaker can plan the vowel sequence (crucially the trigger of deletion), the more likely deletion is to occur. In the following section, we elaborate on these production planning trends and discuss the implications for a listener's capacity to correctly parse the sequence and retrieve the corresponding words.

4.4 Discussion

This study examined hiatus resolution in Madrid Spanish using a forced aligner to automatically categorise /a#e/ and /e#a/ tokens. In so doing, the objective was to test a directional asymmetry underlying the Production Planning Hypothesis: factors that increase the ease of production planning should have a greater effect when the target precedes the trigger than when the trigger precedes the target because the trigger must be planned early enough to apply the process to the target. In this section, we elaborate on our findings. We first discuss how our overall results for hiatus resolution in Madrid Spanish compare to descriptions found in the literature for Spanish and what this suggests about Madrid Spanish phonology (section 4.4.1). We then turn to the Production Planning Hypothesis in section 4.4.2, discussing what the results indicate for how

speakers apply phonological processes like cross-word deletion. In section 4.4.3, we turn to intelligibility as an alternative motivation for variability in vowel deletion, and consider how listeners may recover the intended lexical items in spite of deletion.

4.4.1 Madrid Spanish hiatus resolution across words

In the current study, we focused on the application of vowel deletion to resolve cross-word hiatus in Madrid Spanish. In contrast to previous work (e.g. Aguillar, 1999; Baković, 2006; Garrido, 2008, 2013), the rate of deletion is relatively high in our study. The high rate of deletion that we observe is likely due to the nature of our data; we analyse casual spontaneous speech between interlocutors who were already familiar with each other. Aguillar (1999) observes that deletion is less common in read speech, which is the main source of data for existing work. Furthermore, recall that we did not extract any tokens with pauses between the two words, which means that we examined only tokens of /e#a/ and /a#e/ where deletion is assumed to be possible. As such, we cautiously refrain from concluding that Madrid Spanish speakers delete vowels in hiatus more often than speakers of other dialects.

Consistent with Garrido's (2008) observations, /a/ is unlikely to undergo deletion in hiatus, at least adjacent to /e/. One topic not examined in this paper is the motivation – or motivations – for /a/ to be retained so strongly over /e/ in Spanish, which could be for the speaker's or the listener's benefit. On one hand, /a/ is articulatorily more extreme, which means that it would require more articulatory effort by the speaker and therefore might be disfavoured for preservation. On the other hand, it is more salient to listeners due to its higher sonority; and /e/ being both the most frequent phoneme in spoken Spanish (Moreno Sandoval et al., 2008) and the epenthetic vowel in Spanish (Eddington, 2001) may mean that the default vowel for perceptual reinsertion or for reconstruction would be [e]. Future work could tease apart these explanations by including pairings of /a/ with other vowels, like /o/, for comparison with /e/, as this vowel has greater information content than /e/ but similar sonority.

Our interest in the current study predominantly relates to process directionality rather than segment position, but our results still speak to positional effects. We infer a preference to delete word-final vowels over word-initial ones, which is consistent with both typological observations

on preserving content at the beginning of words (e.g. Barnes, 2006 and Keating, 2006) and models of speech recognition (e.g. Frauenfelder and Komisarjevsky, 1987; Goldinger et al., 1996). We further discuss the effect of a vowel's quality and position in the sequence in the following section, beginning with the implications of our results for the Production Planning Hypothesis (Wagner, 2011, 2012; Tanner et al., 2015; Kilbourn-Ceron, Wagner and Clayards, 2016; MacKenzie, 2016; Kilbourn-Ceron, 2017; Tamminga, 2018; Bailey, 2019).

4.4.2 Production planning as a source of variation

The primary goal of the current study was to test a directional asymmetry implicit to the Production Planning Hypothesis, which proposes that cross-word phonological processes are inherently variable because a speaker will not always be able plan their triggering context. The increased ease of lexical retrieval found for frequent words (e.g. Segui et al., 1982) therefore partly explains the higher rates of reduction observed in these words (e.g. Jurafsky, 1996). However, this common explanation for reduction does not directly motivate processes occurring in adjacent words; it would be surprising for the reduction in word 1 to be conditioned by the frequency of word 2 under the conventional hypothesis of frequency-based reduction. The PPH instead considers processes and their environments, rather than solely the targeted segments and words, giving it the possibility of explaining such phenomena. We therefore reframe frequency and probability effects as being a question of planning contexts for process application, considering both the segments undergoing a phonological process (targets) and those that are part of the environment required for the alternation (triggers). More specifically for hiatus resolution in Madrid Spanish, the process is the deletion of /e/ (the target) when an adjacent /a/ (the trigger) has been planned. We predicted that /e/ deletion would be more affected by factors capturing the ease of production planning in /e#a/ than in /a#e/ because in /a#e/ the trigger (/a/) is consistently planned before the target (/e/), while the same is not true for /e#a/ sequences.

Our results are consistent with the Production Planning Hypothesis in two ways. First, our results demonstrate that deletion of vowels in hiatus is sensitive to lexical frequency and conditional probability such that the ease of planning the sequence is associated with deletion rates. The finding that greater lexical frequency increases deletion in the other word suggests that an account like the PPH is required, rather than this variation being explained through frequent

words being reduced more often (e.g. Jespersen, 1922; Zipf, 1929; Lindblom, 1983; Aylett, 2000; Bard et al., 2000; Bush, 2001; Jurafsky et al., 2001; Bell et al., 2003). Second, our results suggest an asymmetry such that deletion of the first vowel in hiatus is subject to stronger effects of lexical frequency and, to a lesser extent, conditional probability. This is further supported by preliminary exploration of the rare cases favouring /a/ deletion; when /a/ deletes in V1 position, the conditional probability effect is not observably different from deletion of /e/ in the same position, and the lexical frequency effects confirm /a/ deletion in V1 being favoured by production planning factors. Deletion of V2 is less subject to production planning effects overall, which is consistent with processes in which the target follows the trigger not being as susceptible to delays in planning. In the following section, we discuss how speakers may modulate deletion to improve intelligibility.

4.4.3 Vowel deletion and speech intelligibility

An alternative explanation for frequency and predictability effects is that speakers attempt to ensure speech intelligibility. Deletion poses problems for parsing. As reflected in models of speech recognition, it is more difficult for the listener to parse the signal when content is missing (e.g. Pisoni et al., 1985; Astheimer and Sanders, 2011; Mitterer, 2011). These models further reflect that it is more challenging to recover missing or altered content at the beginning of the word compared to at the end of the word, as confirmed by psycholinguistic experiments (for discussion, see Frauenfelder and Komisarjevsky, 1987; Goldinger et al., 1996; Kul, 2007).[47] Based on this, the preference to delete word-final vowels rather than word-initial vowels in cross-word hiatus resolution benefits the listener by being less difficult to parse. Furthermore, deletion is likely to be less problematic for parsing words that are highly frequent or that are especially likely in a given context (i.e. words with high conditional probability), as demonstrated by Astheimer and Sanders (2011). On one hand, this does not directly suggest that speakers intentionally modulate the level of reduction to ensure that their interlocutors can

[47] An apparent counterexample to this edge asymmetry is found in Lavie (2007). Lavie's proposal for speech recognition is based on exemplar models of words being stored in the lexicon. Under this interpretation, the task of parsing a word is not rendered more difficult due to vowel deletion; the variant with deletion is expected to already be stored and is therefore accessible during lexical retrieval. However, this may not apply to words that have not previously been heard with vowel deletion – in this case, a listener would presumably be required to use top-down information or infer the possibility of deletion by comparing the token with deletion to other exemplar clouds.

recover the intended meanings. On the other hand, speakers may modulate reduction, or it may be that variation in reduction is not actively suppressed by speakers because deletion correlates with recoverability. We expect that speakers do not directly modulate deletion to benefit listeners, but that the factors that increase the ability to plan phonological contexts also happen to be those that facilitate lexical retrieval for the listener. Otherwise, if speakers were actively considering the intelligibility of their utterances, we expect that deletion could "erroneously" occur in low-frequency and low-predictability contexts because speakers would plan the process, but fail to plan whether that process should be inhibited. We leave the question of intentional speaker modulation for future work.

Finally, another possibility exists to facilitate recovery of the deleted vowel or improve intelligibility of the word. Specifically, our results are based on a single measure of vowel deletion: whether the closest match for the surface sequence is V1, V2 or V1V2, as determined by the aligner. During preliminary data exploration, we confirmed that F1 and F2 values across the sequence were consistent with the aligner's categorisation. However, we leave for future work the question of whether acoustic cues were available to convey that vowel deletion had occurred, and what the quality of the deleted vowel was. For instance, coarticulatory effects on nearby vowels and consonants may be suggestive of the underlying vowel sequence. Additionally, deletion with compensatory lengthening (e.g. /a#e/ → [a:]) and assimilation (e.g. /a#e/ → [aa]) are both collapsed into deletion (e.g. /a#e/ → [a]) in the current analysis. The two former processes could provide listeners with additional cues that the underlying sequence contains two vowels, for example significantly longer duration or the presence of two peaks in the intensity contour.

4.5 Conclusion

The goal of this study was to use cross-word hiatus resolution in Spanish as a test case for directional effects in production planning, as expected following the Production Planning Hypothesis. We found a strong preference for /a/ to be preserved regardless of its position in the sequence, alongside evidence that planning-based factors are associated with deletion rates in cross-word vowel-vowel sequences. Crucially, we observe two generalisations of particular interest. First, our results suggested a directional asymmetry such that the effects of production

planning are not identical when the triggers of a process precede the targets compared to when the targets precede the triggers. Second, our results show that deletion in a word can be significantly affected by the frequency and predictability of an adjacent word, which is not expected if the results can be attributed to the reduction of frequent or predictable words. Based on this, we conclude that hiatus resolution patterns in Madrid Spanish support the PPH, rather than alternative explanations such as the reduction of frequent or probable words.

The method used for this study is still relatively new. Automated classification using a forced aligner with rules for optional vowel deletion provided a large number of tokens that could be mined for more fine-grained effects. While a lot can be gained from this technique in phonetic and phonological research, it is worth highlighting some of its limitations. Most notably, the aligner makes categorical decisions with minimum segment durations of 30 milliseconds, which means that it cannot capture smaller levels of reduction (like gliding instead of deletion) unless it has been trained on suitable data. Additionally, the aligner does not indicate whether any cues were used to compensate for the deletion -- for example, the aligner does not code that the vowel that was not deleted is a bit longer or has a higher amplitude than it would otherwise have had. Despite the potential limitations of automated classification, however, we have found that this method serves its purpose effectively. In particular, it categorises data in a holistic way, and can enable analyses that may otherwise be infeasible due to difficulties with finding individuals to manually code tokens (e.g. when accessing native speaker coders is challenging, or when manual coding is not feasible due to perceptual similarity between categories or due to the corpus size).

References

Astheimer, L. B., and L. D. Sanders. (2011). Predictability affects early perceptual processing of word onsets in continuous speech. *Neuropsychologia* 49 (12): 3512-3516.

Aylett, M. P. (2000). *Stochastic suprasegmentals: Relationships between redundancy, prosodic structure and care of articulation in spontaneous speech*. PhD diss., University of Edinburgh.

Bailey, G. (2019). Ki(ng) in the north: Effects of duration, boundary, and pause on post-nasal [g]-presence. *Journal of the Association for Laboratory Phonology* 10(1): 3. 1–26.

Baković, E. (2006). Hiatus resolution and incomplete identity. In: Fernando Martínez-Gil and Sonia Colina, (eds.), *Optimality-Theoretic Studies in Spanish Phonology*: 62–73.

Bard, E. G., A. H. Anderson, C. Sotillo, M. P. Aylett, G. Doherty-Sneddon, and A. Newlands. (2000). Controlling the intelligibility of referring expressions in dialogue. *Journal of Memory and Language* 42:1-22.

Bard, E. G., and M. P. Aylett. (1999). The dissociation of deaccenting, givenness, and syntactic role in spontaneous speech. In Proceedings *of the 14th International Congress of Phonetic Sciences,* ed. by J. J. Ohala, Y. Hasegawa, M. Ohala, D. Granville, and A. C. Bailey, 1753-1756. Berkeley: University of California, Berkeley.

Barnes, J. (2006). *Strength and Weakness at the Interface: Positional Neutralization in Phonetics and Phonology: Volume 10 of Phonology and phonetics*. Berlin: Walter de Gruyter, 292 pp.

Baroni, A. (2014). *The Invariant in Phonology: The role of salience and predictability*. Doctoral thesis submitted to Università degli Studi di Padova.

Bayley, R. (1995). Consonant cluster reduction in Chicano English. *Language Variation and Change* 6:303-326.

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

Bell, A., D. Jurafsky, E. Folser-Lussier, C. Girand, M. Gregory, and D. Gildea. (2003). Effect of disfluencies, predictability, and utterance position on word form variation in English conversation. *Journal of the Acoustic Society of America* 113:1001-1024.

Billerey-Mosier, R. (2000). Lexical effects on the phonetic realization of English segments. *UCLA Working Papers in Phonetics* 100.

Blevins, J. (2005). The Role of Phonological Predictability in Sound Change: Privileged Reduction in Oceanic Reduplicated Substrings. *Oceanic Linguistics*, Volume 44, Number 2, December 2005: 517-526.

Boersma, P. and D. Weenink. (2018). *Praat: doing phonetics by computer* [Computer program]. Version 6.0.17, retrieved from http://www.praat.org/.

Borrelli, D. (2002). Raddoppiamento Sintattico in Italian: A Synchronic and Diachronic Cross-Dialectical Study (*Outstanding Dissertations in Linguistics*). Routledge.

Bush, N. (2001). Frequency effects and word-boundary palatalization in English. In *Frequency and the emergence of linguistic structure*, ed. by Joan L. Bybee and Paul Hopper, 255ñ80. Amsterdam: John Benjamins.

Clopper, C. G., and J. B. Pierrehumbert. (2008). Effects of semantic predictability and regional dialect on vowel space reduction. *The Journal of the Acoustical Society of America*, 124(3), 1682–1688.

Connine C, D. Blasko, and D. Titone. (1993). Do the beginnings of spoken words have a special status in auditory word recognition? *J Mem Lang*. 32 (2):193–210.

Casali, R. F. (1996). Resolving Hiatus. Dissertation submitted to the University of California in Los Angeles. 148 pages.

Eddington, D. (2001). Spanish epenthesis: formal and performance perspectives. *Studies in the Linguistic Sciences* 31. 33–53.

Frauenfelder, U. H. and L. Komisarjevsky Tyler. (1987). The process of spoken word recognition: An introduction. *Cognition* 25: 1-20.

Gahl, S., Y. Yao and K. Johnson. (2012). Why reduce? Phonological neighborhood density and phonetic reduction in spontaneous speech. *Journal of Memory and Language* 66 (4): 789-806.

Glowacka, D. (2001). Unstressed vowel deletion and new consonant clusters in English. *Poznan Studies in Contemporary Linguistics* 37: 71-94.

Goldinger, S.D., D. B. Pisoni, and P. A. Luce. (1996). Speech perception and spoken word recognition: Research and theory. In N.J. Lass (Ed.), *Principles of Experimental Phonetics*. St. Louis: Mosby. 277-327.

Gregory, M. L., W. D. Raymond, A. Bell, E. Fosler-Lussier, and D. Jurafsky. (1999). Effects of collocational strength and contextual predictability in lexical production. *Proceedings of the Chicago Linguistic Society* 35: 151-166.

Guirao, M. and M. García Jurado. (1990). Frequency of Occurence of Phonemes in American Spanish. *La phonétique* Volume 19, numéro 2: 135-149.

Guy, G. (1977). A new look at –t, -d deletion. In Fasold, R. W. and R. W. Shuy (eds) *Studies in Linguistic Variation*. Washington DC: Georgetown University Press, 1-11.

Harris, L. N., and C. A. Perfetti. (2016). Lexical Stress and Linguistic Predictability Influence Proofreading Behavior. *Frontiers in Psychology* 7 (96): 1-11.

Hume, E. (2004). Deconstructing Markedness: A Predictability-Based Approach. Proceedings of the Annual Meeting, Berkeley Linguistics Society 13: 182-198.

Jespersen, O. (1922). *Language: its nature, development and origin*. George Allan and Unwin.

Jones, Z., M. M. Dailey, and C. G. Clopper. (2016). Semantic predictability mediates effects of phonetic reduction and social variation on word recognition. *The Journal of the Acoustical Society of America* 140: 3215.

Jurafsky, D., A. Bell, M. Gregory, and W. D. Raymond. (2001). Probabilistic relations between words: Evidence from reduction in lexical production. In *Frequency and the emergence of linguistic structure*, ed. by Joan L. Bybee and Paul Hopper, 229ñ54. Amsterdam: John Benjamins.

Katz, L., L. Brancazio, J. Irwin, S. Katz, J. Magnuson, and D. H. Whalen. (2012). What lexical decision and naming tell us about reading. *Reading and Writing*, 25(6): 1259–1282.

Keating, P. A. (2006). Phonetic encoding of prosodic structure. In Harrington, J. and Tabain, M., editors, *Speech Production: Models, Phonetic Processes and Techniques*, Macquarie Monographs in Cognitive Science, chapter 11, pages 167–186. Psychology Press, New York.

Kilbourn-Ceron, O. (2017). *Speech production planning affects variation in external sandhi*. PhD thesis, McGill University.

Kilbourn-Ceron, O., M. Wagner, and M. Clayards. (2017). The effect of production planning locality on external sandhi: A study in /t/. In: *Proceedings of the 52nd Meeting of the Chicago Linguistics Society*, Chicago: Chicago Linguistics Society.

Kul, M. (2007). *The principle of least effort within the hierarchy of linguistic preferences: external evidence from English*. Doctoral thesis submitted to the Universitet Im. Adama Mickiewicza w Poznaniu. 199 pp.

Labov, W. (1972). *Sociolinguistic Patterns*. Philadelphia: University of Pennsylvania Press.

Lavie, R-J. (2007). Exemplar theory in linguistics: a perspective for the cognitive subject. *Proceedings the 11th Congress of Cognitive Linguistics*, Bordeaux, 19-21 May 2005.

Lema, J. (1987). *An examination of vowel epenthesis in Spanish*. Master's thesis submitted to Simon Fraser University.

Lindblom, B. (1983). *Economy of speech gestures. In The production of speech,* ed. by P. F. MacNeilage, Heidelberg: Springer-Verlag: 217-245.

Lindblom, B. (1990). *Explaining phonetic variation: A sketch of the H&H theory. In Speech Production and Speech Modelling,* ed. by W. J. Hardcastle and A. Marchal: 403-439. Dordrecht: Kluwer Academic Publishers.

Lubbers, M., and F. Torreira. (2016). *Praatalign: an interactive Praat plug-in for performing phonetic forced alignment.* Retrieved from https://github.com/dopefishh/praatalign.

MacDonald, M. G. (1989). The influence of Spanish phonology on the English spoken by United States Hispanics. In P. C. Bjarkman and R. M. Hammond, (eds.) *American Spanish Pronunciation*. Washington, DC: Georgetown Press.

MacKenzie, L. (2016). Production planning effects on variable contraction in English. *Penn Working Papers in Linguistics* 22(2): 121–130.

Mallet, G. M. (2008). *La liaison en français : descriptions et analyses dans le corpus PFC*. Thèse de doctorat. Université Paris X.

Marslen-Wilson W, P. Zwitserlood. (1989). Accessing spoken words: the importance of word onsets. *J Exp Psychol*: *Hum Percept Perform* 15 (3): 576–585.

Mitterer, H. (2011). Recognizing reduced forms: Different processing mechanisms for similar reductions. *Journal of Phonetics* 39: 298-303.

Moren Sandoval, A., D. T. Toledano, R. de la Torre, M. Garrote and J. M. Guirao. (2008). Developing a Phonemic and Syllabic Frequency Inventory for Spontaneous Spoken Castilian Spanish and their Comparison to Text-Based Inventories. In: *Proceedings of the International Conference on Language Resources and Evaluation 2008* (LREC 2008) from 26 May - 1 June 2008 in Marrakech, Morocco.

Nelson, N. and A. Wedel. (2017). Phonetic cues can be reduced in response to minimal pair competition. *91st Annual Meeting of the Linguistic Society of America* (Austin, TX).

Moon, S.-J., and B. Lindblom. (1994). Interaction between duration, context, and speaking style in English stressed vowels. *The Journal of the Acoustical Society of America* 96: 40-55.

Pate, J. K., and S. Goldwater. (2011). Predictability effects in adult-directed and infant-directed speech: Does the listener matter? In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*.

Pisoni, D. B., H. C. Nusbaum, P. A. Luce, and L. M. Slowiaczek. (1985). Speech Perception, Word Recognition and the Structure of the Lexicon. *Speech Communication*, 4(1-3): 75–95.

Pluymaekers M., M. Ernestus and R.H. Baayen. (2005). Lexical frequency and acoustic reduction in spoken Dutch. *Journal of the Acoustical Society of America* (Oct.) 118 (4): 2561-2569.

R Development Core Team. (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. http://www.R-project.org.

Salasoo A, D. Pisoni. (1985). Interaction of knowledge sources in spoken word identification. J *Mem Lang*. 24 (2): 210–231.

Santa Ana, O. (1992). Chicano English evidence for the exponential hypothesis: a variable rule pervades lexical phonology. *Language Variation and Change* 4: 275-288.

Schiel F. (1999). Automatic Phonetic Transcription of Non-Prompted Speech, *Proc. of the ICPhS* 1999. San Francisco, August 1999: 607-610.

Schuppler, B., M. Ernestus, O. Scharenborg, and L. Boves. (2011). Acoustic reduction in conversational Dutch: A quantitative analysis based on automatically generated segmental transcriptions. *Journal of Phonetics* 39: 96-109.

Schuppler, B., S. Grill, A. Menrath, J. A. Morales-Cordovilla. (2014). Automatic Phonetic Transcription in Two Steps: Forced Alignment and Burst Detection. *Statistical Language and Speech Processing* 8791: 132-143.

Segalowitz, S., and K. Korri. (2001). Lexical Access of Function versus Content Words. *Brain and language* 75: 376-89.

Segui, J., J. Mehler, U. Frauenfelder, J. Morton. (1982). The word frequency effect and lexical access. *Neuropsychologia* 20 (6): 615-27.

Sotillo, C. F. (1997). *Phonological reduction and intelligibility in task-oriented dialogue*. PhD diss., University of Edinburgh.

Stevens, J. (2011). Vowel duration in second language Spanish vowels: study abroad versus at-home learners. Arizona Working Papers in SLA & Teaching, 18: 77-104.

Subirats, C., and M. Ortega. (2012). *Corpus del Español Actual* <http://spanishfn.org/tools/cea/english>.

Tamminga, M. (2018). Modulation of the following segment effect on English coronal stop deletion by syntactic boundaries. *Glossa: a journal of general linguistics* 3(1): 1–27.

Tanner, J., M. Sonderegger, and M. Wagner. (2017). Production planning and coronal stop deletion in spontaneous speech. *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, 8(1): 15.

Tily, H., S. Gahl, I. Arnon, N. Snider, A. Kothari, and J. Bresnan. (2009). Syntactic probabilities affect pronunciation variation in spontaneous speech. *Language and Cognition*, 1(2).

Torreira, F., and M. Ernestus. (2011). Vowel elision in casual French: the case of vowel /e/ in the word c'était. *Journal of Phonetics* 39: 50-58.

Torreira, F., and M. Ernestus. (2010). The Nijmegen corpus of casual Spanish. In N. Calzolari, K. Choukri, B. Maegaard, J. Mariani, J. Odijk, S. Piperidis, and D. Tapias (Eds.), *Proceedings of the Seventh Conference on International Language Resources and Evaluation* (LREC'10). Paris: European Language Resources Association (ELRA): 2981-2985.

Wagner, M. (2012). Locality in phonology and production planning. *McGill Working Papers in Linguistics* 22(1): 1–18.

Zipf, G. K. (1949) [1972]. *Human behavior and the principle of least effort. An Introduction to human ecology*. New York: Hafner Publishing Company.

Zipf, G. K. (1929). Relative frequency as a determinant of phonetic change. *Harvard Studies in Classical Philology*, 15: 1–95.

Chapter 5

Conclusion

This thesis used a corpus-based approach to examine the role of grammatical structure and cognitive processing in variation. Put together, the three studies highlight that variation is structured and thereby reveals properties of both speech production mechanisms and the grammar. Section 5.1 begins with a brief overview of the studies comprising this thesis. In section 5.2, I turn to a general discussion of the implications of these studies. In particular, the section explores how the results inform our understanding of the relationship between phonological variation on one hand and prosody, morphology and production planning on the other. Section 5.3 concludes the thesis.

5.1 Overview of Studies

Each study in this thesis provides a test case in which probing the structure underlying variation sheds light on motivations for variation. In this section, I offer a brief summary of those studies, following the order of presentation in this thesis; section 5.1.1 summarises Study 1 on Laurentian French prominence (Chapter 2), 5.1.2 recapitulates Study 2 on high-vowel tenseness in Laurentian French (Chapter 3), and 5.1.3 outlines Study 3 on hiatus resolution in Madrid Spanish (Chapter 4). The following section (§5.2) then discusses some of the themes that emerge throughout the thesis, as well as the implications of the results across the studies.

5.1.1 Study 1

As previously discussed, prominence is conventionally described as being assigned to the final syllable of the accentual phrase in French, but previous work has observed that this is not always

the case. In particular, prominence has been found on penultimate syllables across the French-speaking world. Study 1 probes the issue of variable prominence assignment by drawing on read speech from Laurentian French speakers in the *Phonologie du français contemporain* corpus. In this study, I tested the hypothesis that Laurentian French (a) has heavy vowels and weight-bearing codas, and (b) exhibits weight sensitivity that motivates penultimate prominence. Under this hypothesis, both heavy vowels and codas were expected to attract prominence.

The results are consistent with the hypothesis that prominence assignment is sensitive to weight: prominence is more likely to be assigned to penultimate and final syllables that are closed, as well as to penults with a heavy vowel. The results further reveal that prominence is realised through both the attraction of the accentual phrase's high tone and through an increase in rhyme amplitude and duration. Together, these results suggest that Laurentian French exhibits both (a) lexical prominence, computed based on relative syllable weight and signalled through increased amplitude and duration, and (b) phrasal prominence, which is a pitch accent that is variably attracted away from the right edge of the phrase to a lexically prominent penult. Furthermore, these results highlight the importance of examining the broader acoustic profile of prominence; studies should not consider only pitch in their investigations of prominence in French.

5.1.2 Study 2

High-vowel tenseness in Laurentian French has posed challenges for empirical studies because tense and lax allophones are acoustically too similar to be effectively classified using only one or two acoustic cues. The distinction is rendered more complex as a result of the many interacting – and typically optional – phonological processes that the literature proposes are required for an account of high-vowel tenseness in non-final syllables. Non-final syllables contrast with final syllables in this respect: in final syllables, the phonological processes that determine high-vowel tenseness are obligatory and therefore the tense/lax quality of high vowels is fully predictable in this position. Study 2 exploits the difference between final and non-final syllables by training an aligner to distinguish tense and lax high vowels in final syllables. The aligner then automatically classifies tense and lax high vowels in non-final syllables using those final-syllable language

models. This process produces the first large-scale corpus of Laurentian French speech that is annotated for the tense/lax quality in high vowels.

The results provide mixed support for processes proposed in the literature. On one hand, there is evidence for phonological processes triggered by featural restrictions (Non-Local Harmony, Local Iterative Harmony, Pre-Geminate Laxing and Pre-Fricative Tensing) and by morphological structure (Retensing from Resyllabification). On the other hand, the results are not consistent with other processes (notably Laxing Disharmony) and Laxing Harmony applying at early stages of the derivation. Despite the complexity of the grammar that emerges from the results, I argue that variation in high-vowel tenseness is both learnable from input (contra Poliquin, 2006) and is determined primarily by phonology rather than by phonetic pressures for reduction or coarticulation.

5.1.3 Study 3

In Spanish, variable vowel deletion resolves cross-word vowel hiatus, and both V1 and V2 are potential targets for deletion in these sequences. The potential to have V1 trigger V2 deletion (trigger before target) or V2 trigger V1 deletion (target before trigger) makes it an ideal test case for directional asymmetries in the application of phonological processes. Under the Locality of Production Planning Hypothesis (PPH), which states that phonological variability results from the limited window of phonological encoding, it is expected that processes in which the target precedes the trigger will be sensitive to factors associated with the ease of phonological encoding (e.g. lexical frequency, conditional probability), but that these factors will have limited or no effect in cases where the trigger precedes the target. This asymmetry is a consequence of planning. When the target occurs first, delays in planning the trigger may result in the phonological process not applying; however, when the trigger occurs first, the trigger is expected to invariably be planned when determining whether the target can undergo the process under examination. To test this prediction, Study 3 investigates cross-word hiatus resolution in spontaneous Madrid Spanish, where /e/ may undergo deletion when adjacent to /a/. Forced alignment is used to classify /e#a/ and /a#e/ tokens as being realised as V1, V2 or V1V2, and Bayesian mixed-effects multinomial regression is used to analyse deletion patterns.

The results are consistent with the PPH. Deletion of /e/ in /a#e/ shows relatively little effect of factors associated with the ease of phonological encoding. However, deletion of /e/ in /e#a/ is significantly and considerably more sensitive to lexical frequency and conditional probability. Further support for the PPH is found in the observation that lexical frequency predicts deletion not only in the word itself, but also in the adjacent word. The effect therefore cannot be attributed solely to the reduction of frequent words.

5.2 Implications and Future Directions

In this section, I discuss the implications of the studies in this thesis, as well as directions for future work. The implications are grouped into two themes, with each one focusing on how phonological variation sheds light on linguistic systems and cognitive processing. In section 5.2.1, I discuss how prosody not only varies as a consequence of phonology, as demonstrated in Study 1, but additionally how prosody results in phonological variation. In section 5.2.2, I turn to the question of how morphological structure interacts with phonological processes.

5.2.1 Prosody and Phonological Variation

In Study 1 (Chapter 2), the goal was to probe phonological motivations for variation in prosody, and the results showed that phonological weight significantly influenced prominence assignment. However, effects of the relationship between prosody and phonology are not confined to Study 1, nor are the full implications of the results in Study 1 explored. While many open questions remain, I here focus on the debate surrounding the status of the foot in French and on the potential for prominence to trigger multiple phonological processes.

There is debate about whether feet are projected at all in French and, if so, what type of foot that is built. As described in Study 1, French is not conventionally analysed as having lexical stress and, instead, prominence is assigned by phrases. Secondary stress would normally be an effective test for foot structure, but secondary prominence in French is – like main prominence – seemingly associated with phrase structure, rather than feet (e.g. Jun and Fougeron, 2000). Drawing on this observation and on the apparent lack of minimal size requirements for lexical words (defined as one binary foot; McCarhty and Prince, 1985), many authors have concluded

that French does not project a foot at all (Andreassen and Eychenne, 2013; Jun and Fougeron, 2000; Mertens, 1987; Verluyten, 1982). On the other hand, analyses of reduplication, of truncation, of vowel deletion and of other phonological processes have led other authors to propose that French does not lack the foot. However, faced with conflicting evidence from phonological restrictions and processes, authors in the latter group do not agree on the type of foot that is built; some propose that it is an unbounded foot (Dresher and Kaye, 1990, see also Selkirk, 1978) while others propose that it is bounded. In the latter group, some propose that it is a trochee (Bosworth, 2019; Bullock, 1994; Montreuil, 1993; Selkirk, 1978), others propose that it is an iamb (Archibald, 1996; Charette, 1991; Garcia, Goad and Guzzo, 2017; Goad and Prévost, 2011; Paradis and Deshaies, 1991; Phinney, 1980; Scullen, 1997), and others still that both iambs and trochees are built at different edges (Goad and Buckley, 2006). Montreuil (1994) suggested that the type of foot proposed will likely be predicted by the analysis of schwa: if schwa is treated as underlying, then trochees are formed; if schwa is not underlying, then iambs are formed.

Study 1 focused on the attraction of the pitch accent to a syllable with lexical prominence (as is common with word-level stress; Gordon, 2014). Although the study did not directly test footing, its results can perhaps speak to that issue. We focused on the difference between acoustic cues to prominence (amplitude, duration and pitch) in the penultimate and final syllables. However, preliminary analyses sought to predict the value of those acoustic cues for each syllable, similar to how the results are shown in plots in Study 1. In those analyses, I found that heavy syllables are associated with lower values for acoustic cues to prominence in the neighbouring syllable. This observation suggests that the proportional difference in acoustic cues between syllables is modulated, perhaps due to the presence of the foot. These preliminary analyses are not sufficient to conclude whether final open syllables are unfooted or not when non-prominent, however, which would be important in order to distinguish between trochees (with final syllables in dependent position) and iambs (with unfooted final syllables). This distinction is crucial for Laurentian French because preliminary data exploration reveals that the penult is often prominent when both the penultimate and final syllables are light.

Study 2 neither tested nor controlled for syllable prominence as part of its models to predict high-vowel tenseness in non-final syllables. However, its results may help elucidate footing in

French, at least in Laurentian varieties. A comparison of model intercepts suggests that penults in words with two or four syllables show a preference to surface as tense; this is also a trend observed for the initial syllable of words with four (or more) syllables. These syllables all have one thing in common to the exclusion of other syllables in our data: syllables that favour the tense variant of high vowels appear in odd syllables from the right edge of even-parity words, that is, words that permit syllables to be exhaustively parsed into binary feet. While this may seem like limited evidence for iterative footing, the same difference between odd- vs. even-numbered syllables emerges in work on high-vowel deletion in Laurentian French (Garcia, Goad and Guzzo, in prep; see Goad, 2020): even-parity words favour iterative foot construction while odd-parity words do not. Combining corroboration of this syllable number effect with Poliquin's (2006) proposal that high-vowel Laxing Harmony is motivated by prominence both for the trigger (in all types of Harmony) and for the target (in Non-Local Harmony), we find some support for feet being iteratively built in Laurentian French.[48]

Together, studies 1 and 2 tentatively suggest that Laurentian French projects feet, and that footing is likely to be iterative. They also highlight how variation (in prominence and in phonological processes more broadly) provides insight into the structure of the grammar. Furthermore, recent work suggests that prominence may condition multiple phonological processes in Laurentian French. For example, in Lamontagne (2014), acoustic cues to prominence were significant predictors of mid-vowel realisation, and, in light of the results of Study 1, there is evidence that the contrast between mid-high and mid-low vowels is most robust in prominent final syllables. Similarly, in Lamontagne (2015), acoustic cues to prominence were significant predictors of the fronting of back vowels, suggesting that /ɔ/ fronting occurs in prominent syllables, while /u/ and /o/ fronting occurs in non-prominent syllables. It is likely that many more processes show these same hallmarks of being conditioned by prominence, a question I leave for future work. This work should crucially compare phonological processes (e.g. diphthongisation, vowel mergers) across varieties of French to establish the relationship

[48] While claims of iterativity in "Standard" French go against consensus or are not generally included in descriptions of French prominence, it has been proposed that prominence in French is iterative (Tranel, 1994) and therefore that footing likely would be. Furthermore, Jun and Fougeron (2002) find experimental evidence that high tones may be assigned iteratively, which again suggests iterative footing, although they analyse the effect as being conditioned by morphology due to the shape of their stimuli. For at least some of their speakers, it may be that iterativity motivates the results they obtain. Many authors (Garde, 1968; Pasdeloup, 1990; Verluyten, 1984) also observe that antepenults can bear (secondary) prominence when main prominence falls on the final syllable.

between these processes and the dialect's prominence system and phoneme inventory. For instance, it is plausible that fortition processes like diphthongisation reinforce vowel contrasts and help to signal prominence shift by making the location of prominence more perceptually salient.

5.2.2 Morphology and Phonological Variation

While not the main focus in this thesis, morphological interactions in phonological variation emerge as a common theme across studies. Beginning with Study 1, recall that the models for prominence in Laurentian French included a factor to code whether the penult was the base-final syllable or not. It was noted that this factor was required; otherwise, weight effects on prominence would be understated due to the presence or absence of a morphological boundary between the last two syllables, which is partly predicted by syllable structure. The data used in this study were too limited to probe morphological effects in detail. An account of morphological effects should crucially test for differences between derivational and inflectional suffixes. On one hand, derivational suffixes may attract prominence due to their status as morphological heads and their consequential prosodic autonomy (van Oostendorp, 2002). This could lead either to prominence being assigned to the suffix rather than the base, or to prominence being assigned to *both* the base and the suffix, as they may form separate prosodic domains (for discussion, see Hannahs, 1995). On the other hand, inflectional suffixes may be associated with increased prominence in the base, as suggested by lengthening being preserved or enhanced in base-final syllables in Laurentian French (Walker, 1984) and by inflectional suffixes often being invisible to stress assignment and therefore outside of the domain of stress assignment in other languages (van Oostendorp, 2002). Furthermore, the possibility of morphological effects in prominence assignment is attested in French: the assignment of secondary prominence is sensitive to morphological boundaries (Delais-Roussarie, 1995; Jun and Fougeron, 2002; Pasdeloup, 1992), suggesting the need to consider word-internal motivations for prominence assignment. As such, future work should build on our observation that morphological effects also condition main prominence by examining prominence assignment in unaffixed words as well as in words with derivational and inflectional suffixes. When considering the effects of morphological structure, word should additionally test differences arising from the phonological shape of morphemes (see

Bullock, 1995. As will be highlighted in the discussion of Study 2, phonological processes may differ in how they interact with resyllabification across morpheme boundaries.

Turning to Study 2, we were interested in two possible morphological effects on high-vowel tenseness in Laurentian French. The first was Retensing from Resyllabification, in which a high vowel is tense due to the resyllabification of a coda consonant into the onset of the suffix. Our results were consistent with Retensing from Resyllabification, which signals an important difference between morpheme boundaries (resyllabification of a coda into onset position triggers tensing) and word boundaries (*enchaînement*, which is resyllabification of a coda into the onset of the following word, does not trigger tensing). For example, while /i/ undergoes laxing in /viɲ/ *vigne* 'vine' and in /viɲ#a#lyk/ *vigne à Luc* 'Luc's vine' due to closed-syllable laxing, /i/ is tense in /viɲɔbl/ *vignoble* 'vineyard'. The second morphological effect of interest concerned derivational explanations for laxing; Laxing Harmony could motivate laxing a high vowel (e.g. /myzik/ [mʏzɪk] *musique* 'music'), and affixation later in the derivation could tense the trigger for Laxing Harmony (e.g. /myzik+al/ [mʏzikal] *musique* 'music'), rendering Laxing Harmony opaque. To test this, I coded whether a high vowel could have been targeted by Laxing Harmony at an earlier stage of the derivation. Our results did not support Laxing Harmony triggered by base-final syllables, perhaps due to the presence of other motivations for laxing and the potential for high vowels to surface as lax with no apparent trigger. Another possible morphological effect is found based on cursory data exploration: prefixes in which high vowels would lax (e.g. /dis/ *dis-* 'NEG', /syr/ *sur-* 'over-') seem to be more likely to surface as lax even when Retensing from Resyllabification would be expected to occur. The explanation for this result may, in fact, lie in prosody: cross-linguistic work (Wennerstrom, 1993; Peperkamp, 1997) and work on French (Hannahs, 1995) show that prefixes can form prosodic words that are independent from the following base. As such, future work should probe morphological effects on high-vowel laxing to examine whether prefixes constitute independent domains for Closed-Syllable Laxing, thereby patterning like word-final syllables in having laxing irrespective of coda resyllabification.

Finally, Study 3 did not probe factors relating to morphology with the exception of a brief observation about /a/ deletion in word-final position. As discussed in a footnote (footnote 6 in Chapter 4), preliminary work (Lamontagne and Torreira, 2017) shows that /a/ is more likely to

delete when it is the feminine suffix. At first, this may seem surprising; considerable literature has demonstrated that morphemes are unlikely to delete in their entirety in order to ensure that the meaning or function associated with morphemes is recoverable from surface forms (e.g. Akinlabi 1996; Casali, 1996). However, Spanish has agreement, and /a/ is therefore generally redundant because other words will generally also convey gender. Consequently, application rates for a phonological process increase where syntax renders that morpheme more recoverable. The word-final feminine /-a/ should therefore delete more frequently when gender has been marked on a preceding word because the feminine suffix would be planned as part of the earlier word and thus speakers may retrieve the word or suffix faster. A similar effect has been found for /s/ lenition in the same corpus of Madrid Spanish; the plural suffix /s/ undergoes significantly more reduction when an earlier word conveys number due to agreement (Torreira and Ernestus, 2012). Future work should further investigate how morphosyntactic processes like agreement facilitate phonological encoding by increasing morpheme predictability. Following the Production Planning Hypothesis, the resulting ease of planning the window for phonological encoding should be associated with higher application rates of phonological processes.

5.3 Conclusion

The three studies that comprise this thesis probed the question of how grammatical structure and cognitive processes relate to phonological variation. It was found that variation is associated with – and perhaps results directly from – the interaction between different linguistic systems. This is further expanded upon by considering how phonological variation results from limitations imposed by cognitive processing. The tools crucially enable large-scale analysis, which allows researchers to elucidate subtle patterns that require considerable data to detect. As these tools improve, we will increasingly understand the structure that underlies “free” variation and discover the wealth of data to which learners have access in order to build their grammar.

General Bibliography

Akinlabi, A . (1996). Featural affixation. *Journal of Linguistics* 32: 239-289.

Andreassen, H., and J. Eychenne. (2013). The French foot revisited. *Language Sciences* 39: 126-140.

Appelbaum, I. (1996). The lack of invariance problem and the goal of speech perception. *Proceeding of Fourth International Conference on Spoken Language Processing*, vol. 3: 1541-1544.

Archibald, J. (1996). *The Acquisition of Syllable Weight and Foot Type*. Fifth Conference on Laboratory Phonology (Northwestern University, July 1996).

Baker, A., D. Archangeli, and J. Mielke. (2011). Variability in American English s-retraction suggests a solution to the actuation problem. *Language Variation and Change* 23(3): 347–374.

Bang, H.-Y. (2017). *The structure of multiple cues to stop categorization and its implications for sound change*. PhD thesis, McGill University.

Beddor, P., R. Krakow, and L. Goldstein. (1986). Perceptual constraints and phonological change: A study of nasal vowel height. *Phonology 3*: 197-217.

Blumstein, S. E., and K. N. Stevens. (1981). Phonetic features and acoustic invariance in speech. *Cognition, 10*(1-3), 25–32.

Bosworth, Y. (2011). *Weight and Feet in Québécois*. PhD thesis, University of Texas at Austin.

Bosworth, Y. (2019). High vowel distribution and trochaic markedness in Québécois. *The Linguistic Review*, 34(1): 39-82.

Bullock, B. E. (1994). Does the French syllable have weight? In: M. Mazzola, (ed.), *Issues and theory in Romance Linguistics*, Washington, D.C: Georgetown University Press, 3–15.

Bullock, B. E. (1995). Prosodic constraints and morphological alignment in French. *Lingua* 96: 95-117.

Canavan, A., and G. Zipperlen. (1996). *CALLFRIEND Canadian French LDC96S48*. DVD. Philadelphia: Linguistic Data Consortium.

Casali, R. F. (1996). *Resolving Hiatus*. PhD thesis, UCLA.

Charette, M. (1991). *Conditions on phonological government*. Cambridge: Cambridge University Press.

Chodroff, E. (2017). *Structured variation in obstruent production and perception*. PhD thesis, Johns Hopkins University.

Coetzee, A. W. (2009). An integrated grammatical/non-grammatical model of phonological variation. In: Y-.S. Kang, J.-Y. Yoon, H. Yoo, S.-W. Tang, Y.-S. Kang, Y. Jang, C. Kim, K.-A. Kim, and H.-K. Kang, (eds.), *Current issues in linguistic interfaces*, vol. 2, Seoul: Hankook Munhwasa, 267–294.

Coetzee, A. W. and J. Pater. (2011). The place of variation in phonological theory. In: J. Goldsmith, J. Riggle, and A. Yu, (eds.), *The Handbook of Phonological Theory*. 2nd ed., Wiley-Blackwell: 401-434.

Côté, M.-H. (2008). *La syllabation et le relâchement des voyelles hautes en syllabe non-finale en français québécois*. Paper presented at Journées PFC, Tulane University.

Dresher, B. E., and J. D. Kaye. (1990). A computational learning model for metrical phonology ». *Cognition* 34: 137-195.

Fagyal, Z. (2003). La prosodie du français populaire des jeunes: Traits héréditaires et novateurs. *Le français aujourd'hui* 143: 45-55.

Fischer, J. L. (1958). Social influences on the choice of a linguistic variant. *Word* 14(1). 47–56.

Fougeron, C. (1999). Prosodically conditioned articulatory variations: a review. *UCLA Working Papers in Phonetics* 97: 1-74.

Fruehwald, J. (forthcoming). The Study of Variation. To appear in the *Oxford Handbook of the History of Phonology*. https://jofrhwld.github.io/papers/fruehwald_variationchapter_2017-10-14.pdf. [Accessed July 3rd, 2020.]

Garcia, G. D., H. Goad, and N. Brambatti Guzzo. (in prep). Variable high vowel deletion and rhythmic structure in Québec French.

Gimson, A. G. (1969). A Note on the Variability of the Phonemic Components of English Words. *Brno Studies in English* 8: 75-79.

Goad, H. (2020). *Marginal phonological structure: Prosodic constituency that you cannot 'hear' in Québec French*. Biennial Conference on Laboratory Phonology 17.

Goad, H., and M. Buckley. (2006). Prosodic structure in child French. Evidence for foot. *Catalan Journal of Linguistics* 5: 109–142.

Goad, H., and A.-E. Prévost. (2011). *A Test Case for Markedness: The Acquisition of Québec French*. Ms., McGill University.

Goldwater, S. and M. Johnson. (2003). Learning OT constraint rankings using a Maximum Entropy model. *Proceedings of the Workshop on Variation within Optimality Theory*, Stockholm University.

Guy, G. R., and S. Boyd. (1990). The development of a morphological class. *Language Variation and Change* 2: 1-18.

Hannahs, S. J. (1995). *Prosodic Structure and French Morphophonology*. Tübingen: Linguistische Arbeiten.

Hayes, B. and C. Wilson. (2008). A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39: 379-440.

Jun, S.-A., and C. Fougeron. (2000). A phonological model of French intonation. In: Botinis, Antonis., (ed.), *Intonation: Analysis, modelling and technology*, Dordrecht: Kluwer, 209-242.

Kilbourn-Ceron, O. (2017). *Speech production planning affects variation in external sandhi*. PhD thesis, McGill University.

Kilbourn-Ceron, O., M. Wagner, and M. Clayards. (2017). The effect of production planning locality on external sandhi: A study in /t/. In: *Proceedings of the 52nd Meeting of the Chicago Linguistics Society*, Chicago: Chicago Linguistics Society.

Labov, W. (1972). Some principles of linguistic methodology. *Language in Society* 1(1): 97–120.

Lamontagne, J. (2014). Examining opposing processes in Laurentian French vowels. *Proceedings of the Conference of the Canadian Linguistic Association*, Brock University (St. Catharines, Canada).

Lamontagne, J. (2015). A Little Forward in Laurentian French: A Variationist Analysis of Vowel Fronting in Laurentian French. *Proceedings of the Conference of the Canadian Linguistic Association*, University of Ottawa (Ottawa, Canada).

Lamontagne, J., H. Goad, and M. Sonderegger. (2018). *Morphological and phonological motivations for prominence shifts in French.* 92nd Annual Meeting of the Linguistic Society of America in Salt Lake City, Utah (USA).

Lamontagne, J., and F. Torreira. (2017). *Production Planning Effects in Sandhi: A corpus study using automated classification*. New Ways of Analyzing Variation 46, Madison, Wisconsin (USA).

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

MacKenzie, L. (2016). Production Planning Effects on Variable Contraction in English. *University of Pennsylvania Working Papers in Linguistics* 22(2): 121-130.

Mertens, P. (1987). *L'intonation du français. De la description linguistique à lareconnaissance automatique*. Katholieke Universiteit Leuven, doctoral dissertation.

Montreuil, J.-P. (1993). *Prosody, morphology and foot formation in Modern French*. Paper presented at the Linguistic Symposium on Romance Languages, Northern Illinois University.

Montreuil, J.-P.. (2004a). Fragmenting weight in Scottish English. In: M. Pulki, (ed.), *La Tribune Internationale des Langues Vivantes* 36, Paris: 114–122.

Montreuil, J.-P. (2004b). *The Computation of weight in English and in Québec French*. First PAC Workshop (April 23-24, 2004, Université de Toulouse le Mirail).

Montreuil, J.-P. (2005). *Le poids phonologique en français québécois*. Ms., University of Texas, Austin.

Oostendorp, M. van. (2002). The phonological and morphological status of the prosodic word adjunct. *Linguistische Berichte Sonderheft* 11: 209-236.

Paradis, C., and D. Deshaies. (1991). Rules of stress assignment in Québec French: Evidence from perceptual data. *Language Variation and Change* 2: 135-154.

Peperkamp, S. (1997). *Prosodic words.* The Hague: Holland Academic Graphics.

Phinney, M. (1980). Evidence for a Rhythm Rule in Québec French. *Proceedings of the Tenth Meeting of the North Eastern Linguistic Society, Cahiers linguistiques d'Ottawa* 9: 369-382.

Poplack, S. 1989. The Care and Handling of a Mega-Corpus: The Ottawa- Hull French Project. In: R. W. Fasold, and D. Schiffrin, (eds), *Language Change and Variation*, Amsterdam and Philadelphia: Benjamins, 411-451.

Sankoff, D., G. Sankoff, S. Laberge, and T. Marjorie. (1976). Methodes d'echantillonnage et utilisation de 1'ordinateur dans l'étude de la variation grammaticale. *Cahiers de linguistique de l'Université du Québec* 6, 85-125.

Scullen, M. E. (1997). *French prosodic morphology: A unified account.* Bloomington, Indiana: Indiana University Linguistics Club Publications.

Selkirk, E. (1978). The French foot: on the status of 'mute' e. *Studies in French Linguistics* 1: 141–150.

Tamminga, M. (2018). Modulation of the following segment effect on English coronal stop deletion by syntactic boundaries. *Glossa: A Journal of General Linguistics*, *3*(1): 86.

Tanner, J., M. Sonderegger, and M. Wagner. (2017). Production planning and coronal stop deletion in spontaneous speech. *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, 8(1): 15.

Thibault, P., and D. Vincent. (1990). *Un corpus de français parlé. Montréal 84: historique, méthodes et perspectives de recherche*. Québec: Département de langues et linguistique, Université Laval.

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

Tranel, B. (1994). *The Sounds of French: an Introduction*. Cambridge University Press, Cambridge.

Verluyten, S. P. (1982). *Investigation on French Prosodics and Metrics*. PhD thesis, Universiteit Antwerpen.

Vincent, D., M. Laforest, and M. Guylaine. (1995). Le corpus de Montréal 1995: Adaptation de la methode d'enquete sociolinguistique pour l'analyse conversationnelle. *Dialangue* 6: 29-46.

Wagner, M. (2012). Locality in phonology and production planning. *McGill Working Papers in Linguistics* 22(1): 1–18.

Weinreich, U., W. Labov, and M. I. Herzog. (1968). *Empirical Foundations for a Theory of Language Change*. Austin: University of Texas Press.

Wennerstrom, A. (1993). Focus on the Prefix: Evidence for Word-Internal Prosodic Words. *Phonology* 10(2): 309-324.

Wheeldon, L. (2012). Producing spoken sentences: The scope of incremental planning. In: S. Fuchs, M. Weirich, D. Pape, and P. Perrier, (eds.), *Speech planning and dynamics: Speech Production and Perception*, vol. 1, Peter Lang: 97-118.

Wilson, C. (2006). Learning phonology with substantive bias: an experimental and computational study of velar palatalization. *Cognitive Science* 30(5): 945-982.

Wiltshire, C. (1999). *Variation in Spanish aspiration and prosodic boundary constraints*. 29th Linguistic Symposium on Romance Languages (University of Michigan).