Title

Sociophonetically-based phonology: An Optimality Theoretic account of /s/ lenition in Salvadoran Spanish

Author
Brogan, Franny Diane

Publication Date
2018

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Sociophonetically-based phonology:

An Optimality Theoretic account of /s/ lenition in Salvadoran Spanish

A dissertation submitted in partial satisfaction of the

requirements for the degree Doctor of Philosophy

in Hispanic Languages and Literatures

by

Franny Diane Brogan

2018

ABSTRACT OF THE DISSERTATION

Sociophonetically-based phonology:

An Optimality Theoretic account of /s/ lenition in Salvadoran Spanish

by

Franny Diane Brogan

Doctor of Philosophy in Hispanic Languages and Literatures

University of California, Los Angeles, 2018

Professor Antonio C. Quícoli, Co-Chair

Professor Norma Mendoza-Denton, Co-Chair

This dissertation examines onset and coda /s/ lenition in the Spanish of El Salvador, a dialect in which this phenomenon is particularly advanced. That is, Salvadoran /s/ weakening is not only pervasive in both syllabic positions but manifests as allophones beyond the traditional tripartite conception. In addition to providing the first sociophonetic account of /s/ in El Salvador, this dissertation aims to address another gap in the literature: while /s/ weakening is the most-studied phonological variable in the history of the field, many accounts fail to make connections between observed patterns and important aspects of phonological theory. In my analysis I show that these connections, which are highly reliant on phonetic principles, are crucial to a more complete understanding of the phenomenon at hand.

Speakers for this study are 72 Salvadorans who participated in sociolinguistic interviews in El Salvador in 2015. I segmented and acoustically analyzed 200 occurrences of phonological /s/ per participant (n = 14,400 tokens) in *Praat* (Boersma & Weenink, 2016), and each token was subsequently coded for linguistic and extralinguistic variables including *allophone*, which was categorized as follows: [s]: a voiceless strident fricative; [z]: a voiced strident fricative; [s^{θ}]: a voiceless approximant resulting from gestural undershoot; [h]: a voiceless glottal fricative; [ɦ] a voiced glottal fricative; and [∅]: deletion of the segment in its entirety.

Using these 14,400 tokens, this dissertation develops a phonetically-based phonological analysis of /s/ lenition within Optimality Theory (Prince & Smolensky, 1993/2004) in which the need to reduce articulatory effort cost (Markedness) while preserving important perceptual distinctions (Faithfulness) drives variation. I model the Salvadoran data via a maximum entropy algorithm, which assigns weights to OT-style constraints instead of strictly ranking them. Maximum entropy not only allows for variation but maximizes it, assigning probabilities to all possible output candidates according to shares of total harmony.

Markedness constraints in the grammar are styled following Kirchner (1998, 2004) and incentivize the weakening of difficult articulatory gestures in phonological environments that exacerbate their biomechanical effort cost. I find, for example, that /s/ is most likely to lenite when its flanking segments are more open (such as two vowels that are [-high]) or disagree with it in coronality, voicing, or both. I find that all five non-deleted variants can be modeled according to the interaction between their composite articulatory gestures and the phonological contexts in which they occur. Faithfulness constraints in the grammar, on the other hand, drive the preservation of perceptual cues in more salient prosodic positions. I find that the impetus to preserve both features [+strident] and [-voice] is highest in the strongest prosodic positions (phrase-initially > word-initially > syllable-initially) and decreases gradually in turn with position strength. Furthermore, I find that these differences are modulated by syllable stress, with tonic syllables blocking lenition at higher rates than atonic syllables.

With respect to demographic factors, this dissertation shows that while language-internal factors establish basic constraints in the grammar, the relative importance of preserving perceptual cues varies for different social groups; this is implemented in the grammar by scaling Faithfulness constraint weights up or down. I find that speakers from San Miguel, rural speakers, older Salvadorans, women, and those with lower levels of education prioritize Faithfulness (i.e., the preservation of important perceptual cues) less than their respective counterparts, resulting in higher rates of effort-based lenition. Furthermore, I find that speakers from these groups not only lenite /s/ at higher rates but are also more likely to both produce particularly nonstandard variants such as [s^{θ}] and lenite /s/ more often and more extremely in more salient prosodic positions.

In sum, this dissertation contributes valuable data about a pervasive phenomenon in an understudied dialect of Spanish, including an in-depth exploration of the social factors that condition it. More broadly, this study presents a nuanced approach to Spanish /s/ lenition that is able to account for onset and coda /s/ weakening within a single analysis by situating the phenomenon within well-established theories of phonetically-based phonology.

The dissertation of Franny Diane Brogan is approved.

Kie Ross Zuraw

Ji Young Kim

Antonio C. Quícoli, Committee Co-Chair

Norma Mendoza-Denton, Committee Co-Chair

University of California, Los Angeles

2018

Dedication

I would like to dedicate this dissertation to the Salvadorans who generously shared their time, experiences, and /s/s with me. I was welcomed into churches, community centers, schools, and even homes with a warmth and curiosity I never could have anticipated. This ambitious dissertation would not have been possible without the kindness my 72 participants, not to mention countless others who helped bring my fieldwork dreams to fruition. I would like to especially thank Payton Phillips Quintanilla, Alvaro Eduardo Quintanilla, Alfaro Alfredo Quintanilla, Aída Guadalupe Gómez de Quintanilla, Aída Gabriela Quintanilla, Christina Quintanilla, and Eduardo Chavez for giving me the best home away from home and making sure I was safe and happy during my time in El Salvador.

I would also like to thank the various professors at UCLA who provided me with exemplary training and unwavering support over the past six years: the late Dr. Claudia Parodi, my original adviser, whose contagious fascination with Salvadoran Spanish planted the seed that blossomed into this dissertation; my committee co-chairs, Drs. Carlos Quícoli and Norma Mendoza-Denton, who continually challenged me to think outside the box; to Dr. Ji Young Kim, who joined my committee once the project was well underway and still provided me with invaluable support and feedback; and Dr. Kie Zuraw, who spent countless hours with me in her office helping to iron out the details of an analysis that was ambitious to say the least. I'm also eternally grateful to the UCLA Linguistics Department faculty, who welcomed me into their classes with open arms and so generously shared their expertise with me.

Thank you also to Statistical Consulting Group (special shout-out to Andy, who is an expert in all things R) for providing the kind of specialized support I would expect from a world-class university like UCLA. I (literally) couldn't have done it without you!

Finally, I am so incredibly thankful to my family and friends for being the best cheerleaders and support systems I every could have asked for. To Mariška Bolyanatz: thank you for helping me fall in love with /s/, and for being my favorite work buddy/gal pal/co-author/superwoman. To my dear friends from Piedmont, Davis, Los Angeles, and beyond: thank you for making me feel so capable each and every day. To my parents, Jeff and Paula Brogan: thank you for raising me to be intellectually curious, for instilling in me a work ethic that has never let me down, and for being my best friends. I've made it this far because of you! And finally, to Charlie: thank you for being my partner, my biggest cheerleader, and for never letting me forget how proud of me you are, even when I'm being totally insufferable. If we can survive this dissertation, we can survive anything!

Table of Contents

List of Figures

List of Tables

Acknowledgements

Funding for fieldwork and data collection was furnished by the UCLA Department of Spanish and Portuguese Ben and Rue Pine Travel Grant, awarded in 2015. Additional research funding was provided by the Department of Spanish and Portuguese Shirley Arora Fellowship and the UCLA Dissertation Year Fellowship (DYF).

Vita

University of California, Los Angeles
M.A., Hispanic Languages and Literatures, 2014

University of California, Davis
B.A., Spanish and Comparative Literature, 2010
Magna cum laude

Universidad de Costa Rica, San Jose
University of California Education Abroad Program, 2009

Chapter 1. Introduction

Sibilant weakening in Spanish is a chief phonological characteristic of dialects such as those of Andalusia, Central America, the Caribbean, and the Southern Cone, and has traditionally been described by researchers within a tripartite system, rendering three possible allophones for a given phonological /s/: [s], representing retention of the sibilant; [h], representing retention of glottal spreading but elimination of oral constriction; and [Ø], representing full deletion. Because /s/ weakening is most pervasive in syllable-final position, the vast majority of studies have focused exclusively on coda weakening. However, this phenomenon also occurs in the syllable onset in some dialects of Spanish. This dissertation examines onset and coda /s/ weakening in the Spanish of El Salvador, a dialect in which this phenomenon is particularly advanced (Lipski, 2000). Salvadoran Spanish additionally proves an excellent vehicle for analysis because of the variation it displays beyond a traditional tripartite conception.

Before presenting the research questions and hypotheses that will guide the present study, it is important to first situate the reader within the dialectological landscape, as particular historical circumstances have rendered El Salvador a uniquely interesting object of linguistic study. We begin with a brief discussion of the peninsular origins of American Spanish. Crucially, records from early expeditions to Spanish America that include extensive passenger lists detailing both where in Spain an individual was from as well as where he/she settled in America reveal a strong Andalusian influence during colonization. This is not surprising, as expeditions departed from Andalusian ports such as those of Sevilla, Huelva, and Cádiz (Canfield, 1981). Settlements in Spanish America were first established circa 1550 in the highland areas of present-day Bolivia, Peru, Ecuador, Colombia and Mexico as well as in Guatemala, Costa Rica, Northwest Argentina, and Andean Venezuela. The next settlements were established approximately 100 years later (circa

1650) in present-day Paraguay, New Mexico, Western and *Porteño* Argentina, El Salvador, Honduras, Nicaragua, Southern Chile, and Uruguay; with the exception of Paraguay and New Mexico, these regions are coastal (Canfield, 1981). During the second half of the colonial period, a number of phonological changes were taking place in Andalusia. As settlers continued to flock to Spanish America, they arrived in port cities and settled nearby in recently-colonized coastal regions. These colonizers brought with them currently-developing dialectal changes; however, these changes did not penetrate the less accessible highland regions, effectively establishing the *tierras altas/tierras bajas* (highland/lowland) divide that we utilize today in Spanish dialectology (Canfield, 1981; Hualde, 2005; Lipski, 1995; among others). For this reason, linguistic innovation in Latin America is thought to be strongly linked to areas with colonial seaports.

Particularly innovative among the *tierras bajas* dialects are those of Honduras and El Salvador, whose "historical circumstances [have] resulted in considerable linguistic variation within a small geographical area" (Lipski, 1994, p. 254). El Salvador in particular was linguistically isolated for the duration of the colonial period, during which no newspapers were published and the importation of books was discouraged. Additionally, there was limited access to formal education, rendering the Salvadoran population "disproportionately weighted in the direction of the rural, illiterate speaker" (Lipski, 1994, p. 255). Consequently, at the time of independence, El Salvador had a miniscule number of literate citizens, and public use of language by those in government and religious positions diverged substantially from the educated 'standards' of Spain and Latin America. Until recently, formal education remained out of reach for many Salvadorans and illiteracy rates remained high; as a result, there has been little opportunity to develop a *norma culta* (an educated urban standard), paving the way for vast sociolinguistic diversity.

Despite this diversity, El Salvador has suffered a disproportionate lack of attention within Hispanic linguistics. While its canonical characteristics were first established in the mid-20th century by Canfield (1953, 1960), little empirical work has been done since to develop a more nuanced understanding of Salvadoran Spanish's phonological tendencies. John Lipski has inarguably contributed the most to our knowledge of the dialect, with papers such as *On the weakening of /s/ in Latin American Spanish* (1984) and *El español que se habla en El Salvador y su importancia para la dialectología Hispanoamericana* ('The Spanish spoken in El Salvador and its importance for Hispano-American dialectology') (2000). Lipski's work on El Salvador has focused primarily on the most salient phonological tendencies of Salvadoran Spanish, including /s/ weakening. Such accounts have been crucial for augmenting our general knowledge of the dialect and have drawn attention to social factors that may condition sociolinguistic variation within this community. However, Lipski himself admits that a more complete understanding of sociophonological variation will require a larger body of linguistic research on this dialect.

Most recently, in a linguistic atlas of Central America edited by Miguel Angel Quesada Pacheco, Azcúnaga López (2010) provides an overview of phonological variation in El Salvador based on controlled elicitations from 80 speakers. In his chapter, Azcúnaga López reports the allophonic variation of each standard phoneme (both vocalic and consonantal) according to the percentage of speakers that produced each variant and in which linguistic contexts. In addition to providing descriptive statistics, Azcúnaga López discusses regional tendencies and compiles them in a dialectal map of El Salvador in which he divides and subsequently characterizes the central-western zone and the north-eastern zone based on observed tendencies of his participants. Nonstandard variation of /s/ documented by Azcúnaga López in his atlas is fairly inconsistent with the findings of this dissertation, likely due to significant differences in methodology. However, the assignment of phonological tendencies to specific geographical regions is extremely helpful in understanding the dialectological makeup of a country in which very little linguistic research has been conducted.

While the work of scholars such as Canfield, Lipski, and Azcúnaga López has laid an important foundation for the present study, advances in the field as well as methodological concerns necessitate further research on Salvadoran Spanish. For example, Canfield and Lipski's traditional dialectological work on El Salvador is largely descriptive. While qualitative studies are essential for informing research questions and hypotheses, a large-scale quantitative study such as this one allows us to document systematic phonological patterns among diverse speakers of this dialect, enabling us to forge a better understanding of the rich variation that exists with respect to /s/ weakening in El Salvador. With this knowledge, we are able to better situate Salvadoran Spanish within the greater dialectological landscape of the Spanish-speaking world.

Furthermore, studies in traditional dialectology such as those of Canfield, Lipski and Azcúnaga López have relied exclusively on impressionistic transcription. That is, the researcher listens to a sound produced by a speaker and transcribes what he/she hears based on his/her linguistic expertise. While impressionistic transcription is an important tool for linguists, we as human listeners do not always perceive speech signals accurately because our knowledge and/or expectations of a language may skew our perception of the sounds we hear.[1] With modern acoustic phonetic software such as *Praat* (Boersma & Weenink, 2016), we can utilize empirical acoustic information in tandem with impressionistic transcription to ensure the precision, accuracy, and consistency of our analyses.

[1] A number of experimental studies have argued that segments that are weakened or deleted are reconstructed during speech perception by means of phonetic detail in the input (see Zimmerer & Reetz (2014) for a complete literature review).

Finally, beyond regional generalizations, no study to date has examined the role of social factors in conditioning Salvadoran /s/ lenition. While Canfield (1953) and Lipski (1994) *do* posit that variation in El Salvador occurs along urban/rural lines, this difference is yet to be shown empirically. Furthermore, given the vast body of research on this sociophonetic variable, an investigation of other social effects such as those pertaining to gender, age, level of education, and socioeconomic status is paramount.

In addition to providing a long-overdue sociophonetic account of /s/ in El Salvador, this dissertation aims to address another gap in the literature: while /s/ weakening is the most-studied phonological variable in the history of the field (Esther Brown and Torres Cacoullos, 2002), many accounts fail to make connections between observed patterns and important aspects of phonological theory. In this dissertation, I argue that these connections, which are highly reliant on phonetic principles, are crucial to a more complete understanding of the phenomenon at hand. With these goals in mind, this dissertation is motivated by the following research questions:

1. What is the nature of the /s/ lenition variable in El Salvador? How often is /s/ lenited and how is allophonic variation distributed?
2. How do language-internal factors such as surrounding phonological environment and prosodic position work together to condition patterns of Salvadoran /s/ lenition, and how can these effects be modeled within a theory of Optimality Theory (OT) (Prince & Smolensky, 1993/2004)?
3. How can theories of effort-based consonant lenition within OT provide a unifying framework to account for patterns of both onset and coda /s/ weakening in El Salvador?

4. How are lenition rates and patterns conditioned by the extralinguistic variables of region, urban/rural origin, age, gender, and level of education? What do these findings reveal about the status of the /s/ weakening variable in El Salvador?

In order to address these research questions, I traveled to El Salvador during August and September of 2015 to record the speech of 72 Salvadorans balanced for region (Santa Ana, San Salvador, San Miguel), urban/rural origin, age (18-40; 41+), and gender. Speakers participated in a sociolinguistic interview lasting 45-60 minutes; the purpose of the sociolinguistic interview was, as always, twofold, as I aimed to forge a more complete understanding of the social, political, economic, and cultural forces at work in El Salvador that might be affecting linguistic change and variation while simultaneously obtaining naturalistic speech samples from my participants. All participants were recorded using an Olympus LS-14 Linear PCM recorder digitized at 44.1 kHz and a 16-bit sampling rate with an attached Audiotechnica ATR 3350 lapel microphone.

Beginning 20 minutes into each recording, the waveform and spectrogram of the first 200 occurrences of phonological /s/ (n = 14,400 tokens) were segmented in *Praat* (Boersma & Weenink, 2016) according to segmentation criteria adapted from Erker (2012, p. 49). Tokens of /s/ were coded for the following social and linguistic factors: region, rural/urban origin, age, sex, level of education, prosodic position, preceding segment, following segment, syllable stress, and stress of the following syllable. Tokens were also coded for *perceived allophone*, which was determined using acoustic criteria from Johnson (2012) in tandem with impressionistic transcription; possibilities included the following:

1. [s]: a voiceless strident fricative
2. [z]: a voiced strident fricative
3. [s^{θ}]: a voiceless approximant

4. [h]: a voiceless glottal fricative
5. [ɦ]: a voiced strident fricative
6. [∅]: a deleted token

The phonetic and phonological characteristics of these variants are discussed at length in Chapters 3 and 4.

Using these 14,400 tokens of Salvadoran /s/ as my dataset, I develop a phonetically-based phonological analysis in which the need to reduce articulatory effort cost (Markedness) while preserving important perceptual distinctions (Faithfulness) drives variation (see Chapter 3 for definitions of Markedness and Faithfulness within an Optimality Theoretic Framework). This analysis hinges on the fact that strident consonants like /s/ require precise and prolonged constriction via agonist force in order to achieve their characteristic turbulence. Kirchner (2004, p. 326) explains, "In sibilants in particular, it is known that such antagonism is required: namely, the tongue-blade constriction is partially opposed by a stiffening and bracing of the sides of the tongue against the molar gumline, to produce a grooved channel for the airflow." In addition to challenges posed at the place of articulation, /s/ is also voiceless, requiring sustained glottal abduction throughout the segment's duration. Given the difficulties associated with producing /s/, it is hardly surprising that Salvadoran Spanish employs so many strategies to mitigate its effort cost. These strategies include voicing (/s/ → [z]), gestural undershoot (/s/ → [s^{θ}]), elimination of the place of articulation (/s/ → [h]), elimination of the place of articulation plus voicing (/s/ → [ɦ]) and deletion (/s/ → [∅]).

In lieu of traditional statistical analysis, I choose to model the Salvadoran data via a maximum entropy algorithm, which assigns weights to OT-style constraints instead of strictly ranking them. Maxent not only allows for variation but maximizes it, assigning probabilities to all possible output candidates according to shares of total harmony. This maximum entropy approach is particularly suited to modeling Salvadoran /s/ lenition as even the most unlikely outputs—such as deletion in word-initial, intervallic contexts—are observed in the data, albeit very infrequently. Furthermore, like a human language learner, maxent generalizes patterns from the data it observes and is then able to make predictions about novel inputs. While more labor-intensive than traditional statistical analysis, an important benefit of modeling data in maxent is the ability to couch the analysis within phonological theory, as it is able to implement an OT-style analysis for data with free variation.

In this dissertation, I find that patterns of Salvadoran /s/ lenition are well-accounted for by the factors of phonological environment, prosodic position, and syllable stress, and that these patterns are then modulated depending on the socio-demographic factors of region, origin, age, gender, and level of education. Specifically, I show that while language-internal factors establish basic constraints in the grammar, the relative importance of preserving perceptual cues varies for different social groups; this is implemented in the grammar by scaling Faithfulness constraint weights up or down.

Markedness constraints in the grammar are styled following Kirchner (1998, 2004) and incentivize the weakening of difficult articulatory gestures in phonological environments that exacerbate their biomechanical effort cost. I find, for example, that /s/ is most likely to lenite when its flanking segments are more open (such as two vowels that are [-high]) or disagree with it in coronality, voicing, or both. These contexts increase the biomechanical effort cost of producing the voiceless strident fricative—either by forcing the active articulator to move a greater distance or by necessitating additional gestures in the case of a following consonant that is [-coronal] and/or [+voice]—and therefore increase the impetus to "be lazy" and produce one of five less effortful variants. I find that all five non-deleted variants can be modeled according to the interaction between their composite articulatory gestures and the phonological contexts in which they occur.

Faithfulness constraints in the grammar, also styled following Kirchner, drive the preservation of perceptual cues in more salient prosodic positions. I find that the impetus to preserve both features [+strident] and [-voice] is highest in the strongest prosodic positions (phrase-initially > word-initially > syllable-initially) and decreases gradually in turn with position strength. Furthermore, I find that these differences are modulated by syllable stress, with tonic syllables blocking lenition at higher rates than atonic syllables. I find similar patterns for constraints that necessitate the preservation of the cues [+coronal] and [+segmental], suggesting that asymmetries in onset and coda weakening can be attributed, at least in part, to the need to preserve important perceptual distinctions in more salient prosodic positions.

With respect to social factors, this dissertation finds that region, urban/rural origin, age, gender, and level of education are all important in determining lenition patterns in El Salvador. Speakers from San Miguel, rural speakers, older Salvadorans, women, and those with lower levels of education prioritize Faithfulness (i.e., the preservation of important perceptual cues) less than their respective counterparts, resulting in an elevated importance of Markedness and, consequently, higher rates of effort-based lenition. Furthermore, I find that speakers from these groups not only lenite /s/ at higher rates but are also more likely to both produce particularly nonstandard variants such as [s^{θ}] and lenite /s/ in more often and more extremely in salient prosodic positions.

This dissertation makes important contributions to the fields of dialectology, sociolinguistics, and phonological theory. To the field of dialectology, it contributes valuable information about a pervasive phenomenon in an understudied dialect of Spanish. To the field of sociolinguistics, it contributes an in-depth exploration of the social factors that condition Salvadoran /s/ weakening, including some striking patterns that represent a departure from canonical theory. Finally, this dissertation contributes to Spanish phonology and phonological theory by presenting a nuanced approach to Spanish /s/ lenition that is able to account for onset and coda /s/ weakening within a single analysis couched within well-established theories of phonetically-based phonology.

This dissertation is organized as follows. In Chapter 2, I provide a comprehensive literature review of Spanish /s/, highlighting the most important findings both in El Salvador and in other /s/ weakening dialects of Spanish. In Chapter 3, I present a series of frameworks for motivating, interpreting, and modeling Spanish /s/ lenition within phonetic and phonological theory, focusing specifically on the nuances of Salvadoran /s/ lenition and the merits of some theories over others in accounting for observed patterns of variation. Chapter 4 presents a detailed account of Salvadoran [s^{θ}], which I argue has been misinterpreted as a solely historical artifact and should be considered as a product of effort-based lenition within the purview of this dissertation. Chapter 5 details the present study's methodology, from data collection to acoustic analysis to building the maximum entropy grammar. Chapter 6 presents the results of this dissertation, beginning with the grammar-centric model and then examining differences according to social groups. Finally, Chapter 7 provides the reader with crucial interpretation and analysis of the data, allowing him or her to make important connections with previous literature and to better understand the implications of the results presented in Chapter 6. Chapter 7 also provides a comprehensive summary of this dissertation and directions for future research on Salvadoran /s/.

Chapter 2. Spanish /s/ lenition

2.1 Introduction

Sibilant weakening in Spanish is a chief phonological characteristic of dialects such as those of Andalusia, Central America, the Caribbean, and the Southern Cone. Beyond well-established geographical differences, sociolinguists have produced a prolific body of research identifying various language-internal and social factors that further explain patterns of /s/ weakening (e.g., Alba, 1982; Carvalho, 2006; Cedergren, 1973; Cepeda, 1990; Chappell, 2013; Dohotaru, 1998; Lipski, 1984; Lipski, 1986b; Lynch, 2009; Terrell, 1977; Valdivieso & Magaña, 1991). While these studies differ in the details of their methodologies and analyses, they are united by the assumptions they make in describing the phenomenon in question in categorical terms. That is, /s/ weakening involves debuccalization and/or deletion, traditionally rendering three perceived allophones for a given phonological /s/: [s], representing retention of the sibilant; [h], representing retention of glottal spreading but elimination of oral constriction; and [Ø], representing full deletion. This chapter presents a thorough review of the traditional /s/ weakening literature as well as more recent studies that document /s/ weakening using continuous acoustic measures. A discussion of the vast body of literature on Spanish /s/ lenition is crucial to understanding the forces at work in Salvadoran Spanish, which is grossly underrepresented in the field.

This chapter begins with a brief history of Spanish /s/, which I believe is fundamental to understanding both geographic and sociophonetic variation of the sibilant. This is particularly important with respect to Salvadoran Spanish, which not only debuccalizes and deletes /s/ but employs additional allophones that will be briefly discussed in this chapter and elaborated on in Chapters 3 and 4. Next, I review the /s/ weakening literature, beginning with geographical (dialectal) variation and then moving into more nuanced discussions of the language-internal and social factors that condition this phenomenon. I then present a brief discussion on the use of instrumental methods in studies of /s/ lenition and conclude with theories on its pathway of diffusion. This last section touches on an important debate about how linguists should conceive of /s/ weakening—both synchronically and diachronically—and the challenges of characterizing such a pervasive phenomenon in a uniform way.

2.2 A brief history of Spanish /s/

The articulation of Spanish /s/ is variable cross-dialectally. The story of this variation is thought to begin with phonological changes that occurred in pre-colonial Spain in which the four-sibilant system of Old Spanish was reduced to a two-phoneme distinction. Figure 1 below illustrates this change.

	Old Spanish	Standard Peninsular Spanish	Gloss
I.	/s/	/s/	
	e.g., /pasa/	/pasa/	‘s/he passes’
II.	/z/		
	e.g., /kaza/	/kasa/	‘house’
III.	/ts/	/θ/	
	e.g., /bratso/	/braθo/	‘arm’
IV.	/dz/		
	e.g., /adzer/	/aθer/	‘to do/make’

Figure 1. Various historical changes affecting sibilants reduce a four-phoneme system to two, adapted from Hualde (2011, p. 99).

The /s/ of Standard Peninsular Spanish that resulted from the merging of /s/ and /z/ has an apicoalveolar articulation in which the edges of the tongue rest against the gums and the tip of the tongue creates tight constriction at the alveolar ridge. Figure 2 below shows this convex articulation.

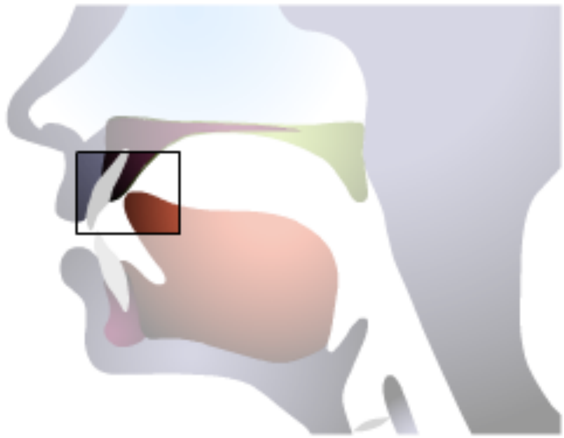

Figure 2. From the University of Iowa Phonetics. Articulation of apicoalveolar /s/ used in Standard Peninsular Spanish ("Sounds of Speech," n.d.).

In opposition to the apicoalveolar /s/ phoneme was interdental /θ/, a voiceless fricative articulated by placing the tip of the tongue between the edges of the front teeth without completely blocking airflow out of the mouth. Figure 3 below illustrates this articulation.

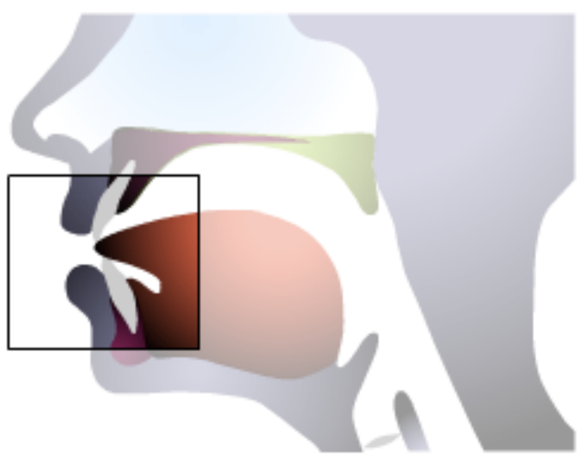

Figure 3. Articulation of interdental /θ/ used in Standard Peninsular Spanish from the University of Iowa Phonetics ("Sounds of Speech," n.d.).

This two-phoneme system, which is often thought of as distinguishing between orthographic 's' (corresponding to /s/) and orthographic 'c' and 'z' (corresponding to /θ/), is referred to by linguists as *distinción* because it allows for minimal pairs such as *casa* /kasa/ ('house') vs. *caza* /kaθa/ ('hunt'). *Distinción* is presently considered to be a feature of Standard Peninsular Spanish and has been categorically maintained in Madrid and parts of northern Spain.

However, by the 15th century, this two-phoneme system had already undergone major transformations in the Spain's southern region of Andalusia, where the two-phoneme distinction had been lost in favor of convergence: orthographic 's', 'c' and 'z' all corresponded to a single phoneme, /s/, formed with a predorsal-dental or -alveolar articulation instead of an apicoalveolar one. As is illustrated in Figure 4 below, this sound is articulated by placing the predorsal region of the tongue against the upper teeth, leaving a small opening through which air escapes from the vocal tract. The tip of the tongue comes down to rest against the lower teeth, creating a concave configuration that can be contrasted with the convex configuration seen in Figure 3 above.

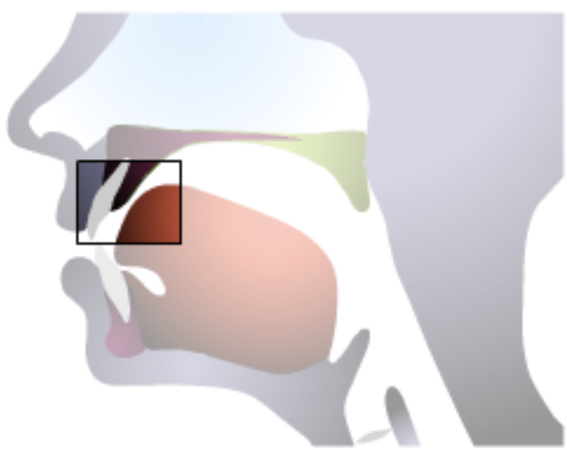

Figure 4. Articulation of predorsal-dental or -alveolar /s/, originally observed in Andalusia from the University of Iowa Phonetics ("Sounds of Speech," n.d.).

Linguists refer to this system, in which predorsal-dental /s/ is used exclusively, as *seseo.*

To complicate matters even more, another change was simultaneously taking place in Granada and other parts of southeastern Andalusia, where some scholars claim the two-phoneme distinction had also been lost in favor of convergence but in the opposite direction: orthographic ‘s’, ‘c’ and ‘z’ were all pronounced as an apicodental fricative perceptually similar[2] to interdental /θ/, referred to as *ceceo.* Linguists have used a variety of symbols to represent this sound, whose articulation varies substantially, including but not limited to: [θ], [s^{θ}], [θ̩], and [θ^{s}]. I will refer to this sound as [s^{θ}] throughout this dissertation.

While traditional linguistic atlases such as those of Navarro Tomas et al. (1933) and Alvar et al. (1976) suggest that each of the three systems (i.e., *distinción, seseo,* and *ceceo*) were used categorically based on geographically-determined dialectal differences, these claims have been widely debated and are, at least in modern-day Spain, considered to be linguistic idealizations. For one thing, as Hualde (2005, p. 154) points out, “Due to recent mobility, nowadays in Andalusian cities such as Seville and Granada it is possible to find speakers who practise *distincción,* speakers who employ *seseo* and speakers who use *ceceo.*” Additionally, and perhaps more importantly, the idea that even a single speaker categorically utilizes one system over another has been widely contradicted; in fact, there appears to be a large amount of intraspeaker variation driven by forces of stigma and prestige. Within Andalusia, “*distinción*, as part of the Peninsular standard, has greater prestige, followed by *seseo*, which can be considered the local standard, whereas *ceceo* is considered ‘rural’ and is evaluated negatively by Andalusian city-dwellers” (Hualde, 2005, p. 154). Indeed, this intraspeaker variation in which speakers utilize all three allophones [s], [θ], and [s^{θ}] in free variation has been attributed to hypercorrection by speakers who aspire to use

[2] Penny (2000, p. 118) describes the articulation of *ceceo* as the “fronting of the tongue body so that the sound acquires some of the acoustic qualities of interdental /θ/.”

prestigious *distinción*. This shift has been empirically documented in traditionally "*ceceante*"[3] Granada (Melguizo Moreno, 2007; Moya Corral & García-Wiedemann, 1995) and Málaga (Ávila Muñoz, 1994; Lasarte Cervantes, 2010; Villena Ponsoda, 1996, 2007) as well as in "*seseante"* Jerez de la Frontera (García-Amaya, 2008), Huelva, and Lepe (Regan, 2017).

Due to the influence of eastern Andalusian colonizers, discussed in Chapter 1, Latin America is thought to be almost entirely *seseante,* employing a predorsal-dental or -alveolar /s/. However, the *ceceo* allophone of eastern Andalusia "appears to be common, but subject to individual variation, in the Central American [country] of El Salvador" (Hualde, 2005, p. 154). The presence of this sound in American Spanish—and specifically in the Salvadoran dialect—will be revisited in Chapter 4, where I offer thorough review of the literature as well as a comprehensive account of the articulatory, acoustic, phonological characteristics of [s^{θ}].

Another *andalusismo* affecting /s/ that is ubiquitous in the *tierras bajas* regions of Latin America is /s/ lenition, which is the focus of this dissertation. Given its prevalence, sibilant weakening has been a topic of interest of Spanish phonologists for some time; the first evidence of the systematic reduction of /s/ was documented in 18th century speech of Andalusians of low socioeconomic status (Lipski, 1984). At this time, /s/ was exclusively vulnerable in syllable-final (coda) position, a marked environment cross-linguistically (Blevins, 1995). While it seems that /s/ weakening was restricted to coda position when this phonological phenomenon was brought to the Americas by Andalusian settlers, over time this *andalucismo* developed independently in different *tierras bajas* regions of Latin America and has become difficult to characterize with a uniform phonological process. Therefore, in an effort to empirically document overarching patterns of Spanish /s/ weakening as well as dialectal nuances, a vast body of sociolinguistic research has

[3] I place both *ceceante* and *seseante* in quotes here to express my skepticism that these regions ever used one system categorically.

emerged in recent years. The following section reviews this literature, beginning with dialectal generalizations and then digging deeper into the language-internal and social factors that have been shown to condition patters of Spanish /s/ lenition.

2.3 Variation in Spanish /s/ lenition

/s/ lenition is the most-studied variable in Hispanic Linguistics, and for good reason: this phenomenon is a rich source of variation conditioned by geographical, linguistic, and social factors. The section reviews the pertinent literature on Spanish /s/ lenition in an effort to situate the research questions, hypotheses, and goals of this dissertation within the current conversation.

2.3.1 Dialectal variation

Spanish is the second-most spoken language in the world, with more than 400 million native speakers across more than 20 countries (Simons & Fennig, 2017). Consequently, dialectologists have documented immense geographical variation, from morphology to syntax to phonology and beyond. While such geographical distinctions can be helpful as we attempt to document the linguistic variation that exists within a single language, "it is important to realize…that the dialects of a language are not discrete entities with sharp boundaries" (Hualde, 2005, p. 18). In the sections that follow, I present geographical variation in patterns of Spanish /s/ lenition, dividing the Spanish-speaking world into regional groups that are helpful for making generalizations. However, it is important to recognize that such generalizations can be limiting; these limitations will be addressed at the end of this section.

2.3.1.1 /s/ lenition in Spain

/s/ lenition is pervasive in the Spanish regions of Andalusia (where it is thought to have originated), Murcia, and Extremadura, as well as in the Canary Islands (Penny, 2000). Additionally, Penny (2000) notes that in recent years this phenomenon has spread north to traditionally /s/-retaining areas of Castile. /s/ weakening in these regions occurs primarily in syllable-final position, both word-medially, as in *casco* [ˈkah.ko] ‘helmet’, and word-finally before pauses, consonants, and vowels, as in *las olas grandes* [laˈho.lahˈgɾan.deh#] (‘the big waves’). While the voiceless glottal fricative [h] and deletion [Ø] are the most frequent outcomes of /s/ lenition in Spain, Gerfen (2001) documents cases of voicing or voicing plus nasal assimilation of [h] to the following consonant, as in *cisne* [ˈsiɦ.ne] or [ˈsiɦ̃.ne] (‘swan’), in Andalusia. Both Penny and Gerfen also document gemination of the following consonant, as in [ˈsi^{h}n.ne], which is uncommon outside of Andalusia.

2.3.1.2 /s/ lenition in Latin America

As I discuss in Chapter 1, there has been a general consensus among dialectologists that /s/ weakening is Latin America is typical of island, coastal, and low-land areas and that its use, therefore, cannot be delineated by national borders. There is a debate as to how this *tierras altas* (highland) vs. *tierras bajas* (low-land) dichotomy came to be. Both Canfield (1981) and Lipski (1994), on one hand, argue that the prevalence of *andalusicmos* in coastal regions of Latin America are the result of historical patterns of migration and colonization. Settlements in Spanish America were first established circa 1550 in the highland areas of present-day Bolivia, Peru, Ecuador, Colombia and Mexico as well as in Guatemala, Costa Rica, Northwest Argentina, and Andean Venezuela. The next settlements were established approximately 100 years later (circa 1650) in present-day Paraguay, New Mexico, Western and *Porteño* Argentina, El Salvador, Honduras, Nicaragua, Southern Chile, and Uruguay; with the exception of Paraguay and New Mexico, these regions are coastal (Canfield, 1981). During the second half of the colonial period, a number of phonological changes were taking place in Andalusia. As settlers continued to flock to Spanish America, they arrived in port cities and settled nearby in recently-colonized coastal regions. These colonizers brought with them currently-developing dialectal changes; however, these changes did not penetrate the less accessible highland regions (Canfield, 1981; Hualde, 2005; Lipski, 1995; among others). Penny (2000) offers an alternate explanation—one that I do not see as mutually exclusive to that proposed by Lipski and Canfield—in which colonial regions that were culturally isolated from the Spanish government, i.e., the later-settled *tierras bajas* regions, were able to be more phonologically innovative because they were free of the forces of linguistic prestige imposed on those who remained within the sphere of influence of Castile.

Recognizing that /s/ lenition has developed somewhat disparately in the numerous low-land regions of Latin America, various typologies have been proposed in an attempt to capture systematic differences in how this phenomenon manifests. Kaisse (1999), for example, uses morphological distinctions to propose a typology for /s/ weakening dialects based on Argentinian and Caribbean Spanish. Her proposal can be seen Figure 5.

All dialects: /s/→ [h]: morpheme-internal and stem-final: ___C			
Type I	Buenos Aires Argentinian	/s/ → [h] prefix-final, __+C word-final, __#C	 [deh.kwiˈðar] *descuidar* 'to neglect' [ˈe.rehˈgran.de] *eres grande* 'you (inf) are big'
Type II	Rio Negro Argentinian & Caribbean II	/s/ → [h] prefix-final, ___+C word-final, ___#{C, V}	 [eˈhal.go] *es algo* 'it's something'
Type III	Caribbean Spanish I	/s/ → [h] prefix-final, __+{C, V} word-final, ___#{C, V}	 [de.haˈser] *deshacer* 'to melt'

Figure 5. Typology of aspirating dialects, adapted from Kaisse (1999).

In Kaisse's proposal, all /s/ weakening dialects do so in morpheme-internal and stem-final (word-internal) positions before a consonant. Dialects in the Type I group extend this rule to apply to prefix-final, preconsonantal environments as well as word-final contexts when the following word begins with a consonant. Dialects in the Type II group weaken /s/ in the same contexts as the previous group and additionally in word-final position before a vowel. Finally, dialects in the Type III group allow weakening in the largest number of contexts, adding prefix-final, prevocalic contexts.

Lipski (1986), in an attempt to compare /s/ weakening patterns across all Latin American dialects, reports the percentage that each allophone [s], [h], [Ø] occurs in five possible contexts: /s/C, /s/#C /s/##, /s/#V́, /s/#V̌. A few generalizations can be drawn from the data. First, that speakers in Bolivia, Costa Rica, highland Colombia and Ecuador, and Guatemala maintain the sibilant almost categorically, except when /s/ occurs word-finally before another consonant (/s/#C). Second, /s/ weakening dialects can be divided into two groups: those that have high rates of reduction in preconsonantal position only (Argentina, Chile, coastal Colombia, Peru, and

Uruguay) and those that weaken /s/ in all five contexts (Cuba, the Dominican Republic, Puerto Rico, Venezuela, Paraguay, coastal Ecuador, El Salvador, Honduras, Nicaragua, and Panama). With a few exceptions, /s/ reduction appears to be most advanced in the Caribbean and Central America, a generalization that is more or less consistent with stereotypes both in the field of Hispanic linguistics and among speakers of these various dialects. However, it is difficult to separate the generalizations presented in Lipski's account from the socioeconomic makeup of these regions versus those in which /s/ weakening is less pervasive. As we will see in the next section, higher-class speakers in /s/ weakening regions are generally less likely to produce nonstandard variants than working class speakers, so it is unsurprising that we find more innovative patterns of weakening in the least-developed regions of the Spanish-speaking world such as those of the Caribbean and Central America.

In an effort to operationalize common vs. innovative patterns of weakening, Penny (2000), proposes a hierarchy of reduction that ranks the relative frequencies of /s/ weakening patterns in Latin American Spanish. Distinct from both Kaisse and Lipski, however, Penny includes cases of syllable-onset /s/ lenition, which is far less common than coda /s/ weakening in Spanish and is salient in very few dialects.

I.	Aspiration of syllable-final /s/ word-internally, word-finally before a word-initial consonant, and phrase finally.	[ˈeh.tah.muˈxe.reh] *estas mujeres* ‘these women’
II.	Aspiration of syllable-final /s/ word-internally and word-finally before a word-initial consonant, with deletion in phrase-final position.	[ˈeh.tah.muˈxe. reø]
III.	Aspiration of syllable-final /s/ word-internally, with deletion word-finally before a word-initial consonant and in phrase final position.	[ˈeh.ta.muˈxe. reø]
IV.	Deletion in all these positions:	[ˈeø.taø.muˈxe.reø]
V.	Aspiration of word-final /s/ before a word-initial vowel with or without deletion of word-internal, syllable-final aspiration.	[ˈe(h).taˈho.ßraø] *estas obras* ‘these works’
VI.	Deletion of word-final /s/ in all positions, including before a word-initial vowel.	[ˈeø.taøˈoßraø]
VII.	Extension of aspiration to word-initial /s/.	[ˈeø.ta.heˈɲo.ra] *esta señora* ‘this lady’

Figure 6. /s/ weakening frequency hierarchy, from the most common patterns of /s/ reduction (I) to the least common (VII), adapted from Penny (2000, p. 150).

According to Penny’s hierarchy, aspiration is more common than deletion, word-final reduction is more common before a consonant than before a vowel, and syllable-onset weakening is the rarest manifestation. It should be noted that there is an ongoing debate about the relationship between coda and onset weakening—that is, is onset /s/ weakening an extension of coda /s/ weakening, or a different process altogether? Various perspectives will be discussed in Section 2.5.

While Penny’s hierarchy does give a nod to weakening in word-initial contexts, it does not account for other cases of syllable-initial reduction, such as that which occurs word-medially as in [ˈka.ha] *casa* ‘house.’ While far less common, Lipski (1985, 1986b) argues that both word-initial and word-medial weakening occur fairly systematically in syllable-onset position in El Salvador and Honduras . Lipski (1985:147) calls his findings evidence of a “phonetic innovation,” stating that "the frequency [of /s/ weakening in onset position] is high enough to be of potential significance as a phonological variable." Indeed, Mason (1994, p. 7) calls the exclusive study of coda /s/ an "oversight," given that onset /s/ weakening is typologically far less common (and therefore more remarkable) than coda /s/ weakening. Onset /s/ weakening has also been attested in dialects of Chihuahua, Mexico (Esther Brown & Torres Cacoullos, 2002; 2003), New Mexico (Esther Brown, 2005), and Cali, Colombia (Earl Brown & Esther Brown, 2012).

While dialectal classifications are helpful for understanding general linguistic patterns in a language with so many diverse speakers, such over-simplifications can be limiting if taken alone. Even within major dialectal areas, one can observe ample geographic and sociolinguistic variation—a reality that is only becoming more pronounced as the flow of people and ideas is facilitated in an ever-more globalized world. With this in mind, the following sections review the literature on language-internal and social factors that have been shown to condition patterns of Spanish /s/ weakening, with a particular focus on those that will be relevant to the analyses presented in this dissertation.

2.3.2 Language-internal factors

In this section, I review the literature on language-internal factors known to affect patterns of /s/ lenition in Spanish. As very little is known about Salvadoran Spanish previous to this dissertation, I focus on findings that have been consistently informative across /s/ weakening dialects. While a plethora of linguistic and usage-based factors have been shown to condition /s/ reduction, I focus here on those that are relevant to this dissertation, whose primary concern is the interaction of constraints that aim to mitigate articulatory effort cost while simultaneously preserving important perceptual distinctions. For that reason, I begin with a discussion of the effects of the surrounding phonological environment (which directly relate to effort cost) and follow with a discussion of factors related to perceptual distinction: prosodic position, stress, and lexical frequency of the host-word. Finally, I provide a brief summary of additional factors that have been shown to condition /s/ weakening but that are not examined in this dissertation.

2.3.2.1 Surrounding phonological environment

In accounts of onset /s/ weakening in Spanish, the preceding segment has consistently been identified as a crucial linguistic predictor. In a study of Chihuahuan Spanish, Ester Brown and Torres Cacoullos (2002; 2003) find that preceding non-high vowels favor weakening significantly more than their counterparts, and that /s/ is retained most frequently after a preceding consonant. Both Esther Brown (2005) and Earl Brown and Esther Brown (2012) find identical results in New Mexico and Cali, Colombia, respectively. These results are unsurprising given that sibilant /s/ requires tight oral constriction and a high tongue body, both of which are more difficult to achieve (i.e., more effortful) when flanked by more open segments, such as non-high vowels. Similarly, these three studies find that the nature of the segment that follows /s/ is a predictor of weakening, although the effect in the Chihuahua data is not statistically significant. As is the case with the preceding segment, a non-high vowel following /s/ favors weakening. Taken together, these results suggest that intervocalic position is the most favorable for onset /s/ weakening, particularly when one or both vowels are [-high].

In traditional studies of coda /s/ weakening, the segment that follows phonological /s/ has consistently proven to be a significant predictor of patterns of sibilant weakening (e.g., Cedergren, 1973; Cid-Hazard, 2003; Lynch, 2009; Ruiz-Sánchez, 2004). However, these studies vary in how they conceive of the 'following segment' factor. Many older studies employ a three-way distinction of the following segment category (consonant/vowel/pause), with generally consistent findings that /s/ is retained most readily before a pause and weakened most readily before a following consonant (e.g., Terrell & Tranel (1978) on Puerto Rico, Cuba, and Venezuela). However, these findings are contradicted by Dohotaru (1998) in Havana, Cuba and Ruiz-Sanchez (2004) in Caracas, Venezuela, who later find that pre-pausal deletion is pervasive, as is word-final deletion regardless of the following segment. One possible explanation for these contradictory findings is that /s/ reduction grew more advanced over time in these innovative Caribbean regions. In Panama, however, Cedergren (1973) finds patterns similar to those later identified by Dohotaru and Ruiz-Sanchez: aspiration is most prevalent pre-consonantally, whereas deletion is preferred before a following vowel or pause.

In more recent accounts of Spanish /s/ lenition, some authors have chosen to expand upon the following-segment distinction and sub-categorize consonants by place and/or manner of articulation. Both File-Muriel (2009) and Valdivieso and Magaña (1991), who categorize following consonants based on manner of articulation, find that following fricatives favor weakening while following stops disfavor it in Colombia and Chile, respectively. Lynch (2009), who sub-categorizes following consonants by both place and manner of articulation, finds that /s/ is weakened most before velar stops and palatal consonants in Miami Cuban Spanish. Other studies yet have chosen to group following segments by features such as [± continuant] (Ruiz-Sánchez, 2004). In a study primarily concerned with deletion vs. sibilant retention and aspiration in Caracas, Ruiz-Sanchez finds that deletion is preferred before consonants that are [-continuant], contradicting the findings of File-Muriel (2009) in coastal Colombia.

While data on the effect of the surrounding phonological environment on /s/ lenition is robust, findings are often left to stand alone, never explained via phonological or phonetic theory.

For example, Lynch finds that rates of /s/ retention are significantly higher preceding particular consonant classes, but the linguistic forces that might be driving this result are never explored. In contrast, this dissertation will conceive of all of its language-internal variables within an effort-based theory of consonant lenition (see Chapter 3), subcategorizing factors such as preceding and following segments purposefully and strategically, so that findings are able to adequately address the theoretical questions at hand.

The following sections review the literature on factors related to perceptual distinction that will be paramount in this dissertation's analysis of Salvadoran /s/ weakening.

2.3.2.2 Prosodic position

Whether /s/ occurs word-initially or word-medially has also been shown to condition rates of weakening. With respect to onset weakening, both Esther Brown and Torres Cacoullos (2002, 2003) and Esther Brown (2005) find higher rates of lenition in word-medial position than in word-initial position. However, Earl Brown and Esther Brown (2012) report contradictory results in Cali, Colombia, where /s/ is reduced to [h] or [Ø] significantly more frequently in word-initial position. With respect to coda weakening, the majority of studies find higher rates of weakening word-finally than word-medially (e.g., Cedergren, 1973; Dohotaru, 1998; Hammond, 1980; Lipski, 1984; Ruiz-Sánchez, 2004; among others).

While prosodic strength relations are seldom if ever mentioned in traditional accounts of /s/ lenition, these results are well-accounted for within such a theory. The basic idea is this: the phonetic realization of an individual speech segment depends, in part, on its position within the larger prosodic structure of the intonational phrase in which it occurs. Keating (2006, p. 169) writes: "The general idea of a strength relation is that prosodic positions are stronger or weaker;

segment/feature phonetic realizations are also stronger or weaker; and segment strength matches position strength, with stronger pronunciations in stronger positions." In this dissertation, which examines both onset and coda /s/ weakening in a single analysis, theories about the phonetic coding of prosodic structure will play an important role in accounting for patterns of lenition in different prosodic positions.

2.3.2.3 Stress

Studies of coda /s/ weakening in dialects such those of as Chile (Lenz, 1940), Miami (Alfarez, 2000; Lynch, 2009), and Cartagena, Colombia (Lafford, 1986) have found that lenition is favored in atonic syllables and is comparatively less frequent in tonic syllables. Both Esther Brown and Torres Cacoullos (2002, 2003) and Earl Brown and Esther Brown (2012) find similar results in their respective accounts of onset /s/ weakening in Chihuahua, Mexico[4] and Cali, Colombia. In this dissertation, the role of syllable stress—like that of prosodic position—will be couched within theories of prosodic strength relations in which more prominent prosodic positions favor stronger segments, which in turn help to maintain perceptual distinctions.

In addition to the stress of the syllable in which /s/ occurs, some accounts of Spanish /s/ weakening have emphasized the importance of surrounding syllable stress. In Chihuahua, Esther Brown and Torres Cacoullos find that /s/ weakens most readily in pretonic position. Contrastingly, in his account of /s/ reduction in Salvadoran and Honduran Spanish, Lipski (1984) shows that /s/ is more frequently maintained when it precedes a stressed syllable regardless of its position within the syllable and word.

[4] In this dialect, the effect was only significant word-initially.

2.3.2.4 Lexical frequency

Cross-linguistically, more advanced phonological reduction has been observed in higher frequency words (Bybee, 2002). This observation holds true with respect to Spanish /s/ lenition, both in accounts of onset /s/ weakening (Esther Brown, 2005; Esther Brown & Torres Cacoullos, 2002, 2003; Earl Brown & Esther Brown, 2012) and coda /s/ weakening (Erker, 2012; File-Muriel, 2009; File-Muriel & Earl Brown, 2010; Minnick Fox, 2006). Such findings can be understood via a systematic production bias, which relies explicitly on the competing needs for the speaker to reduce articulatory effort cost while preserving important perceptual distinctions necessary for communication. As Lindblom (1983) explains, speakers will undershoot articulatory targets to save effort to the extent that communication is not compromised, i.e., target words are still recoverable. In her account of exemplar dynamics, Pierrehumbert (2001, pp. 146-47) applies this idea to lenition processes, explaining that "a systematic lenition bias causes the distribution of exemplars to shift… and causes an increase in variance." Furthermore, "Since the high frequency words are used more often than the low frequency words, their stored exemplar representations show more numerous impacts of the persistent bias toward lenition." In other words, the effect of a word's lexical frequency on lenition rates is gradient, with higher frequency words more recoverable regardless of phonetic variation (in our case, variable /s/ lenition) in their production.

2.3.2.5 Additional factors

Other language-internal factors have been shown to influence /s/ weakening, though their effects sometimes vary cross-dialectally. For good measure, findings with respect to these factors are summarized below but will not be considered in this dissertation, whose goal is to account for these patterns within a theory of effort-based consonant lenition.

Some studies, such as those of Terrell (1979) in Cuba and Cepeda (1995) in Chile, have argued for the importance of word length, showing that /s/ is more resistant to deletion when it appears in monosyllabic words. Offering a possible explanation for this finding, Terrell (1979, p. 609) claims that "the danger of misunderstanding increases" when /s/ is reduced in a shorter word, and therefore "some effort (unconscious) is made to preserve the phonetic substance of monosyllabic words." In a similar vein, Earl Brown and Esther Brown (2012) find that rates of onset /s/ weakening are positively correlated with word length in Cali, Colombia, an observation that might be explained via a theory of polysyllabic shortening (Lehiste, 1972): the durations of syllables in longer words get shorter, necessitating the reduction of individual segments in turn.

Another factor frequently included in accounts of /s/ weakening is the morphological status of /s/. Word-final /s/ occurs as a morpheme in words such as /kome+s/ *comes* ('you eat'), in which it serves as a 2nd person singular verbal marker, or /mesa+s/ *mesas* ('tables') in which it serves as a plural marker, as well as non-morphemically in words such as /martes/ *martes* ('Tuesday'). Studies of /s/ weakening have aimed to test the hypothesis that /s/ will be retained more frequently when it plays a morphological role (as in *comes* or *mesas*) than when it does not (as in *martes*) but have yielded mixed results. While some studies find no relationship between morphological status of /s/ and rates of lenition (e.g., Cepeda, 1995; Hundley, 1987; Ranson, 1992), others find more weakening of inflectional /s/ (e.g., Lafford, 1986; Lynch, 2009), which contradicts the hypothesis in question. Poplack (1980) suggests that the situation is more complex than simply comparing morphological and non-morphological /s/, finding that morphological /s/ is retained when there is no other disambiguating information within the noun phrase (NP) in which the word occurs. In other words, the inflectional /s/ of /mesa+s/ might be deleted in the NP *las mesas* [las ˈme.sa] (‘the tables’) but not in *tiene mesas* [ˈtje.neˈme.sas] (‘he/she/it has tables’), because the article in the former serves to disambiguate the noun’s number.

2.3.3 Social factors

Language-internal factors clearly play an important role in conditioning /s/ weakening, but they cannot be separated from the social characteristics of their speakers. While “social variation in American Spanish has been very much less well studied than geographic variation” (Penny, 2000, p. 161), a number of recent studies have strived to identify the social factors that influence patterns of Spanish /s/ lenition. These studies almost always rely on methods and theories associated with first-wave sociolinguistics, which looks for correlations between linguistic variation and macro-social categories such as age, gender, and socioeconomic status. The following sections review previous findings on the effects of age, gender, rural/urban place of origin, and level of education on patterns of Spanish /s/ lenition.

2.3.3.1 Age

Across different dialects of Spanish, age is sometimes selected as a conditioning factor for /s/ weakening. In these studies, younger speakers tend to weaken /s/ more frequently than older speakers (Alba, 1990; Cedergren, 1973), perhaps because young people tend to be the leaders in linguistic change (Eckert, 1989, 2000). However, if /s/ weakening becomes associated with dispreferred demographics, younger speakers may become less likely to weaken /s/ to avoid stigmatization. For instance, Lynch (2009) finds that young Cuban Spanish speakers in Miami were more likely than their older counterparts to maintain /s/. He hypothesizes that /s/ weakening is associated with more recently-arrived Cuban immigrants from whom these younger speakers prefer to diverge. In the same vein, younger speakers in a border community in Uruguay aspirate /s/ more frequently than their older, [s]- producing counterparts. Carvalho (2006) claims that this is a change from above, or an innovation toward the prestigious [h] of Montevidean speakers.

Contrastingly, Lafford (1986) finds that younger speakers in Colombia use more 'standard' variants than older speakers in reading tasks because they tend to be more educated than the generations that came before them. She argues that higher levels of education and literacy cause speakers to speak more prescriptively, which may be attributed to their awareness of prestigious variants and prescriptive grammar, as well as the cues provided to them by Spanish's transparent orthography that are not enjoyed by illiterate individuals. As a result, the mental grammars of literate and illiterate speakers may contain different underlying forms for words that are frequently subject to /s/ weakening. For this reason, education and, by extension, socioeconomic status (SES), are thought to be highly reliable predictors of sociophonetic variation. With regard to /s/ weakening in Spanish, speakers who pertain to higher socioeconomic classes typically use significantly higher rates of the standard variant (e.g., Carvalho, 2006; Cedergren, 1973; Ruíz-Sánchez, 2004).

2.3.3.2 Gender

Additionally, gender is one of the most commonly-employed demographic categories in sociolinguistics, and for good reason: Labov (2001), after more than three decades of research, claims to have found the cleanest, clearest results of his career in linguistic comparisons across genders. His findings, summarized in a set of principles, state that women use more standard forms with respect to stable variables. Similarly, when linguistic change comes from above, women are quick to adopt the prestigious form. However, when change comes from below—beginning in the speech community—women are quicker to adopt new, innovative forms. This pattern, which is so consistent in the research that it has become a staple in sociolinguistic theory, is highly apparent in studies of /s/ weakening in Spanish. In a number of first-wave studies that have looked at the correlation between gender and /s/ production, women tend to produce the local prestigious form (usually the sibilant, but not always) at significantly higher rates than men (Cedergren, 1973; Cepeda, 1990; Cid-Hazard, 2003; Poplack, 1979; Terrell, 1986). Many scholars have offered interpretations for such patterns, including the ideas that women are oriented more toward overt social norms (Trudgill, 1972) or that men are more immersed in the workforce and consequently exposed to a wider variety of linguistic variation.

2.3.3.3 Rural vs. urban origin

Finally, both in Latin America and Spain, the dialectal division between urban and rural areas has been well established. Sánchez (1983) describes the process by which urban centers provide speakers with more contact with standardized forms of language through mediums such as formal education and technology, leading to the stigmatization of non-standard variation that is more pervasive in countryside dialects. The absence of standardizing effects in rural communities has important implications, as linguistic isolation makes these communities both more resistant to changes toward standard language forms as well as more likely to innovate even further away from the standard (Cotton & Sharp, 1988). This rural-urban division is exemplified in sociolinguistic accounts of Panama in which urban speakers actively try to distinguish themselves from rural speakers by avoiding stigmatized pronunciations resulting from processes such as debuccalization and /r/ lateralization. As a result, Robe (1960) finds high levels of hypercorrection in the urban dialect.

Similarly, studies such as those of Lopez Morales in San Juan, Puerto Rico (1983) and Cedergren in Panama City (1973) show that urban and rural groups both engage in high rates of /s/ weakening yet pattern differently with respect to their preferred variants. Such differences indicate that rural and urban speakers in Latin America, despite living in close geographical proximity, often pertain to different speech communities within their respective dialectal regions: rural speakers are more likely to weaken /s/ to [h], while their urban counterparts are more likely to delete /s/ entirely. In El Salvador, both Canfield (1953) and Lipski (1994) intuit that sociolinguistic variation occurs along urban/rural lines, but this difference has yet to be shown empirically. This dissertation takes rural/urban place of origin as a proxy for socioeconomic status, which is difficult to otherwise operationalize in El Salvador due to the socioeconomic makeup of the country.

2.3.3.4 Education

Finally, level of education has been identified as important predictor of /s/ weakening. This is unsurprising given findings that speakers with less education are more likely to innovate away from a prestigious variant (Labov, 1972, 2001). Speakers with higher levels of education may be more cognizant of (or simply have more exposure to) a linguistic standard. Indeed, studies such as Dohotaru’s (1998) account of coda /s/ weakening in Havana reveal that sibilance retention is highest among the most educated Cubans and deletion is most prevalent among the least educated.

2.3.3.5 Additional factors

In addition to the extralinguistic factors that will be considered in this dissertation, speech style has consistently been identified as a good predictor of /s/ weakening patterns: as speech style decreases in formality, lenition increases. While this inverse relationship appears to exist for speakers of all socioeconomic groups (Fontanella de Weinberg, 1974), Lafford (1986) finds that differences in formal and informal speech are even more pronounced for groups that are particularly attuned to the social stigma associated with weakening. In her study of Colombian Spanish, Lafford compares rates of weakening among various social groups across four speech styles (spontaneous speech, semi-formal speech, paragraph reading, and word list reading) and finds that young speakers weaken more than older speakers in spontaneous speech yet less in the more formal speech styles. Lafford interprets these findings as an indication that, while /s/ weakening continues to be extended in conversational speech styles, young speakers—who experience more social and geographical mobility than older generations and are therefore more aware of the prestige associated with sibilant retention—avoid lenition in careful speech, where they are more conscious of their pronunciation. Because this dissertation is concerned with modeling variation in naturalistic speech and data collection has been designed to produce a uniform speech style, this factor will not be taken into account in its analysis.

Section 2.3 has summarized three important forces in shaping patterns of /s/ lenition in Spanish: geography, language-internal factors, and social factors. While the vast majority of these studies conceive of /s/ weakening within a tripartite, segmental system (i.e., [s], [h], [Ø]), a more recent dialogue among Hispanic linguists has challenged whether such an approach is adequate. The following section reviews the literature on instrumental approaches to /s/ weakening, focusing on the arguments made in favor of these methods over traditional ones, the findings they yield, and the possible shortcomings of such approaches.

2.4 Instrumental studies of /s/ lenition

Arguments against the use of the traditional tripartite system ([s], [h], [Ø]) when documenting patterns of /s/ weakening are numerous and stem from both practical and theoretical considerations. A chief concern, for example, is the difficulty of representing a tripartite phenomenon via statistical analysis. Sociolinguistic studies on /s/ weakening typically model their data using variable rule analysis or binary logistic regression, which explain the relationship between one categorical dependent variable (phonetic classification of /s/) and one or more sociolinguistic factors. However, the requirement that the dependent variable be binary forces the researcher to collapse two or more phonetic categories into one. Consequently, many scholars such as Lynch (2009) choose to collapse the weakened allophones [h] and [Ø] and report effects on only the presence vs. absence of sibilance. Such models have continually posed a challenge for researchers who are forced to treat as binary a variable that is not dichotomous in nature.

Additionally, theoretical critiques of the segmental approach hinge on the fact that /s/ weakening, like other lenition processes, is a gradient phenomenon. In other words, phonetic realizations of /s/ fall on an acoustic continuum, with a canonical sibilant representing one extreme and phonetic zero representing the other. Historically, due to practical limitations before the widespread availability of acoustic analytic tools such as *Praat* (Boersma & Weenink, 2018), non-phonetician linguists had little choice but to impose impressionistic, segmental labels on tokens of /s/. However, variationists within Hispanic linguistics have begun to question the theoretical soundness of modeling a gradient phonological process in a categorical way. In an effort to promote the use of instrumental measures over categorical labels, several recent papers on /s/ weakening have shown that linguistic factors also condition variation along continuous acoustic dimensions (Erker, 2010; File-Muriel & Earl Brown, 2010; Univaso et al., 2014; among others)

Two such measures that have been used both in studies of natural speech and in laboratory research on Spanish /s/ are spectral center of gravity (COG) and duration. COG is a measure of the average frequency of fricative energy and can be calculated by the equation COG= $\sum fI/\sum I$, where *I* is amplitude (in dB) and *f* is frequency (in Hz) (Erker, 2010, p. 13). In other words, this weighted average is calculated by multiplying values for amplitude in a given spectral slice by the frequencies at which they occur, summing those products, and then dividing that value by the total sum of amplitudes. The two key components of this equation, amplitude and frequency, are crucial measures in differentiating fricatives. A tighter constriction results in a higher amplitude of turbulent air passing through a smaller space (Johnson, 2012). Additionally, Johnson (2012, p. 156) notes that the articulation of [s] is characterized by "airflow past the teeth [that] produces periodic vortices in the airflow which contribute high-frequency components to the spectrum of [s]." Within this framework, a fricative's COG should be a good measure of the degree of constriction achieved in its production and has, in fact, been shown to correlate well with articulatory properties (Tabain, 2001).[5]

A reduction in frication duration, a temporal measure, is another potential indicator of weakening. Within an effort-based approach, the articulation of [s] is particularly difficult "due to the action of antagonistic muscles that is required to hold the articulator in place. (Strident consonants like *s* require particular precision to achieve their characteristic strong turbulence.)"

[5] See Chapter 4 for a spectrographic comparison of various phonetic realizations of /s/.

(Zuraw, 2009, p. 8). Without enough time to execute a gesture precise enough to produce strong turbulence, gestural undershoot and subsequent weakening may occur.

Previous studies of /s/ weakening via instrumental measures such as duration and COG have shown that acoustic variation in the data is systematic. Minnick Fox (2006), for example, shows that a number of factors including lexical frequency, word predictability, and phonetic context are good predictors of variation in COG and duration of /s/. In a study of a similar nature, File-Muriel and Earl Brown (2010) utilize measures of duration, COG, and percent voicelessness to show that factors such as local speech rate and stress have a significant effect on these outcome variables.

Of those championing the use of acoustic measures to reflect gradient phenomena, Erker (2010) is the first to explicitly compare the efficacy of segmental and instrumental approaches. In his account of /s/ weakening by Dominican Spanish speakers in New York, Erker argues that segmental approaches obscure systematic patterns of lenition for a number of reasons. First, he argues that the descriptive adequacy of segmental approaches is diminished by both the presence of within-category variation (e.g., two tokens categorized as [s] may be quite dissimilar acoustically) and the fact that temporal and spectral weakening are not perfectly correlated (e.g., tokens labeled as [s] are significantly shorter when they appear before a following consonant as opposed to a vowel, yet there is no significant difference in COG). Erker then runs two types of statistical models on his data. The first is a binary logistic regression following the traditional segmental approach. He also runs two linear regression analyses using continuous measures of frication duration and COG as dependent variables and argues that the continuous analyses account for almost three times more variance than the categorical one when taking identical predictors into account. With these results, Erker argues that the continuous models are a more comprehensive reflection of patterns in the data. He also champions these methods for reducing transcription bias. Furthermore, as both and Erker (2010) and File-Muriel and Earl Brown (2010) discuss, instrumental measures of weakening are not perfectly correlated: different phonological environments may favor or disfavor one dimension of weakening over another, providing information that is "impossible to capture using traditional IPA categories, which collapse all relevant acoustic cues into several categorical labels" (File-Muriel & Earl Brown, 2010, p. 53).

These scholars at the forefront of instrumental approaches to /s/ weakening in Spanish demonstrate some advantages over segmental accounts. However, it is important to acknowledge the possible shortcomings of such analyses, including the difficulty of disentangling articulatory mechanisms of speech production unrelated to lenition—which often have an effect on temporal or spectral characteristics of fricatives—from phonological weakening. Carney and Moll (1971), for instance, find that anticipatory articulation before a following high vowel results in a higher tongue body for /s/, the effects of which can be seen in the acoustic signal. Similarly, in a study of coarticulation effects on fricatives, Tabain (2001) shows that /s/ has a lower spectral peak in the context of rounding by comparing electropalatographic (EPG) recordings of various CV tokens. Tabain's results are consistent with those of Shadle and Scully (1995), who suggest that [u] has a strong acoustic effect on alveolar fricatives because the lip rounding of [u] causes a whistle-like sound source.

Such coarticulatory effects on the acoustic dimensions of /s/ pose problems for the instrumental study of /s/ weakening. While an instrumental approach might reveal that /s/ appears "weaker" before rounded vowels, for example, this systematic variation is not necessarily the result of lenition. In this sense, instrumental approaches to /s/ weakening likely *do* reveal patterns obscured by segmental representations, but these patterns are not always meaningful in the context of phonological weakening. That being said, I feel it is important to acknowledge that the line between "phonetic" and "phonological" is not clear cut. Rather, authors such as Campos-Astorkiza (2014) make a compelling case that phonological processes like assimilation derive from phonetic conditions related to gestural magnitude and timing. Within this framework, it could be argued that coarticulation, i.e., increased gestural overlap resulting from shorter and/or less distinct gestures, is in fact a type of weakening. What I am arguing here, in contrast, is that coarticulatory effects may sometimes affect acoustic measures in ways that are *not* meaningful or informative for phonology. For example, the whistle-like sound produced by a rounded vowel might lower the spectral COG of the /s/ in solo /solo/ ('only'), leading us to mistakenly conclude that this /s/ is quantitatively "weaker" than the /s/ in sala /sala/ ('living room') or silla /siʝa/ ('chair').

Furthermore, while segmental accounts have been criticized for imposing discrete categories on outputs that are inherently continuous, there may be some merit to associating tokens of /s/ with perceptual categories. Previous work has shown that different allophones of [s] have discrete social or indexical meanings for listeners. For instance, Carvalho (2006) demonstrates that aspiration is prestigious in a border Uruguayan community, as it is associated with high-status Montevideo speakers. That is, her study shows that naïve speakers and listeners differentiate among allophones of /s/ and assign indexical meaning to these uses. Similarly, Tassara and Duque (1987, p.323) explore attitudes toward /s/ weakening in Chile, and their results show that "realizations of /s/ are associated with a very conscious evaluation on the part of the community, and that lack of familiarity with the rules that require different phonetic variants of /s/ in each type of situation produces a disqualification of the speaker, and in some cases, heavy social sanctions" (my translation). In other words, segments have social meaning.

In the first instrumental account of /s/ weakening to incorporate social factors (to my knowledge), Brogan and Bolyanatz (to appear) subject Salvadoran /s/ weakening data to a series of analyses—two instrumental analyses taking duration and COG as their dependent variables and a third analysis in which they combine traditional segmental and instrumental approaches. The authors find that, while the instrumental models are able to account for patterns of /s/ weakening based on language internal factors, neither picks up on any social variation despite the segmentally-based model selecting age, gender, and region as highly significant predictors of /s/ weakening. These findings (or lack thereof) confirm suspicions that instrumental analyses alone may not be adequate to document patterns of sociophonetic variation, as they may obscure potential social meaning.

Despite these limitations, Brogan and Bolyanatz caution against throwing the baby out with the bathwater and advocate for a hybrid method that combines the principles of categorical and instrumental analyses. Specifically, they map instrumental measures onto segmental categories, use these measures to calculate how [s]-like each weakened variant of /s/ is, and then arrange those variants on a numerical 'constriction scale' used as the dependent variable in an ordinal logistic regression. The authors conclude that this hybrid approach is optimal in that it is able to capture the gradient nature of consonant weakening while acknowledging the importance of both segmental categories and the acoustic measures that define them. While this dissertation adopts a segmental approach, instrumental measures of duration and COG will be used in Chapter 4's in-depth analysis of Salvadoran [s^{θ}].

In the following section, I conclude my review of the literature on Spanish /s/ lenition with a discussion of the various theories of diffusion that have been proposed with respect to this phenomenon.

2.5 Theories on the diffusion of /s/ lenition

Thus far, this chapter has been concerned with synchronic patterns of /s/ lenition. This section, in contrast, will focus on the diachronic diffusion of /s/ weakening, or how this phenomenon is extended from one context to another over time. I first present various proposals in the ongoing debate about diachronic pathways of /s/ weakening and then discuss how I will reconcile these conflicting perspectives in this dissertation.

Much of the debate about pathways of /s/ weakening boils down to a simple question: are onset weakening and coda weakening one process or two? Ferguson (1990: 64) proposes that cross-linguistic /s/ lenition follows two paths. The first is the Spanish type, in which weakening begins in syllable-final positions before a consonant and then extends "last, if at all, to word-initial position." The second, the Greek type, begins in syllable-onset, intervocalic position and then extends word-initially and, finally, to preconsonantal positions. However, scholars such as Mendez Dosuna (1996), Lipski (1999), and Terrell (1977, 1979) disagree with this model, arguing for the use of a hierarchy similar to Penny's (see Section 2.3.1.2) to account for both synchronic frequency and diachronic change. Similar to what Penny proposes, Lipski argues that weakening begins in preconsonantal contexts and that the most advanced stage of lenition for almost all dialects is weakening in word-final, prevocalic position. In very few dialects, and only in those in which word-final, prevocalic weakening occurs, /s/ weakening may be generalized to word-initial and syllable-initial/word-internal contexts. Lipski (1999, p. 199) notes that this last configuration is "in its incipient phase" in "some marginal varieties of Spanish", referring to—assumedly—the Spanish El Salvador, where onset /s/ weakening is fairly common.

A number of studies have since challenged the framework set forth by Lipski and Mendez Dosuna and generally accepted in the field. In fact, a cross-dialectal comparison of /s/ weakening in seven dialects presented by Earl Brown and Esther Brown (2012) presents a much more complex picture in which two distinct types of /s/ weakening dialects emerge: those that reduce /s/ significantly more in word-final, preconsonantal positions than in prevocalic positions, and those that reduce /s/ at similar (not significantly different) rates in preconsonantal and prevocalic positions and also show evidence of syllable-onset weakening. Earl Brown and Esther Brown (2012, p. 102) write that "there may be something special about the intervocalic phonological context in word-final position, and further that this lack of significant difference could potentially be used as a diagnostic to predict the existence of initial /s/ reduction in varieties of Spanish."

Another interesting conclusion drawn from studies of dialects in which both onset and coda /s/ weakening are salient (i.e., Esther Brown & Torres Cacoullos, 2002, 2003; Esther Brown, 2005; Earl Brown & Esther Brown, 2012) is that onset and coda /s/ weakening represent two distinct variables. Earl Brown and Esther Brown (2012, pp. 105–106) write:

> Individual analyses of initial /s/ tokens and final /s/ tokens revealed different orderings of factor groups (by magnitude of effect) and were suggestive of two separate variables (/s-/ and /-s/) as opposed to one (/s/). Also, quantitative analyses of reduction rates of /s/ tokens in different syllable positions and phonological contexts…do not point to a unified process of /s/ reduction proceeding from preconsonantal contexts to ultimately include word-medial intervocalic contexts.

Faced with the data presented in these accounts, I agree that it is difficult to argue for /s/ diffusion patterns as they are traditionally conceived of. However, keeping in mind that Salvadoran Spanish patterns with this dialectal group, I argue in this dissertation that a in lieu of so-called 'diffusion pathways,' we should think about these emerging dialectal groups as having developed different constraint rankings within an Optimality Theoretic Framework. That is, an OT analysis couched within theories of phonetically-based phonology can account for seemingly divergent patterns of weakening without having to conceive of onset and coda /s/ weakening as unrelated or separate phenomena. Chapter 3 explores this proposal in more depth.

2.6 Conclusion

In this chapter, I have presented a panorama of Spanish /s/ lenition that incorporates historical, dialectological, and sociolinguistic perspectives. Given the lack of empirical research—sociophonetic or otherwise—on the Spanish of El Salvador, this vast body of literature serves as a foundation for investigating patterns of /s/ weakening in Salvadoran Spanish and situating these findings within contemporary phonological and sociolinguistic theory. Indeed, while variation of Spanish /s/ lenition has been thoroughly documented, the literature as a whole fails to offer convincing phonological motivations for the observed patterns. In the next chapter, I review literature on effort-based approaches to consonant lenition and offer a theoretical framework that I argue is able to encompass the various aspects of Spanish /s/ weakening in a single, unifying analysis. Specifically, Chapter 3 will review phonetic motivations for /s/ lenition as well as theories of phonetically-based phonology, a framework within Optimality Theory that hypothesizes that phonological patterns are determined by phonetic knowledge.

Chapter 3. /s/ lenition within phonetic and phonological theory

3.1 Introduction

Serving as the theory-heavy companion to the previous chapter, this chapter presents a series of frameworks for motivating, interpreting, and modeling Spanish /s/ lenition within phonetic and phonological theory. Much of this chapter focuses specifically on the nuances of Salvadoran /s/ lenition and the merits of some theories over others in accounting for observed patterns of variation.

This chapter begins with the basics by reviewing the phonetic and phonological relationships between [s] and [h] and consequently motivating /s/ weakening from a linguistic perspective. This discussion lays the foundation for Kirchner's (2004) effort-based account of consonant lenition, which is presented in Section 3.5 after I summarize the key aspects of two theories on which it is predicated: Optimality Theory (henceforth, OT) (Prince & Smolensky, 1993/2004) and phonetically-based phonology in Sections 3.3 and 3.4, respectively. I then shift gears and turn to my data, exploring how Salvadoran /s/ lenition might be accounted for using Kirchner's effort-based approach. Finally, I address two key challenges in modeling variable phonological phenomena in OT: accounting for free variation in a system that relies on strict constraint ranking, and the role of social factors in a constraint-based grammar.

3.2 The phonetic and phonological relationships between [s] and [h]

In order to understand the phonetic and phonological theory that has been applied to Spanish /s/ lenition, it is first helpful to review the basics of source-filter Theory (Fant, 1960), which is crucial to understanding the articulatory relationship between [s] and [h].

3.2.1 Source-filter theory

Speech is typically conceived of as the result of a two-stage process within the vocal tract. First, sound is generated in larynx—the *source*. This sound consists of a spectrum of acoustic energy (i.e., a fundamental frequency and harmonics), which is then *filtered* by the resonant properties of the vocal tract according to its configuration. In other words, the filter modulates harmonic amplitudes produced by the source differently depending on its shape, effectively allowing some frequencies to pass through while blocking or attenuating others. Figure 7 illustrates this basic idea.

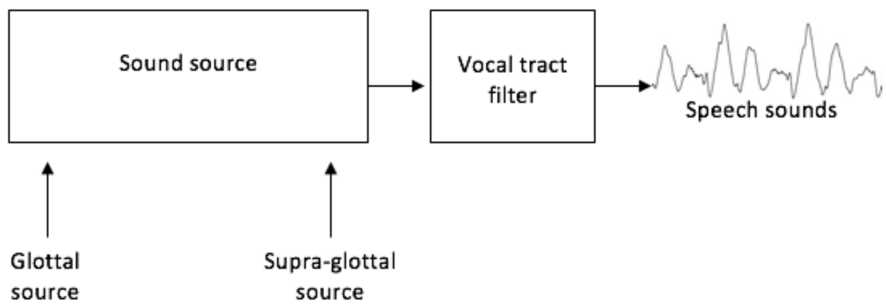

Figure 7. The source of a speech sound originates in the larynx and is then filtered depending on the configuration of the vocal tract, producing a unique speech sound.

Speech output is, simply put, the result of the combined effects of the source and filter. Consequently, if the laryngeal source remains constant but the vocal tract configuration changes, the resulting speech sound will be different. When the harmonics of the sound generated in the source move into the vocal tract, the size and shape of the airway results in harmonic resonance in which some harmonics are made stronger via sympathetic vibration while others become weaker. No new harmonics are added or subtracted, but they come out with either more or less amplitude than when they entered depending on the vocal tract's configuration.

Voiceless fricative sounds begin with turbulent noise in the supra glottal source. As noise generated in the larynx moves up through the vocal tract, additional noise is generated at the place of articulation, where turbulent airflow collides with articulators as it attempts to escape the vocal tract. In the case of [s], air flows through the tightly constricted space between the tongue and the alveolar ridge and then bounces off the upper teeth, causing displacement and high amplitude noise characteristic of sibilants. If the constriction in the filter is relaxed and/or the place of articulation shifts, the amplitude of the noise generated in the vocal tract decreases, resulting in a non-sibilant approximant (see Chapter 4 for a thorough account of this process in Salvadoran Spanish). If the filter is reconfigured so that the noise in the vocal tract is eliminated entirely (i.e., there is no constriction of articulators in the vocal tract and therefore no place of articulation), we are left with only the source, and the resulting sound is one produced simply by noise generated in the larynx, i.e., [h]. This process by which a speech sound is reduced to a laryngeal consonant is commonly referred to as *debuccalization.*

Within the context of this chapter, which proposes a theoretical framework for analyzing Spanish /s/ lenition based on a myriad of phonetic and phonological theory, it is helpful to think of the /s/ → [h] change as one in which the 'filter' component of 'source-filter' is lost. The question that will be addressed in this chapter is, then: why would a language want to eliminate the filter function if it plays a crucial role in differentiating speech sounds? The following section presents a standard treatment of debuccalization within phonological theory, which speaks to the phonological relationship between [s] and [h]. Taken together, an examination of both the phonetic and phonological relationships between [s] and [h] provides a foundation for investigating the articulatory motivations and perceptual implications of debuccalization, i.e., the costs and benefits of eliminating a sound's place of articulation, which will be the focus of later sections of this chapter.

3.2.2 Debuccalization in Autosegmental Theory

Debuccalization occurs cross-linguistically and is one of a handful processes considered within the purview of consonant lenition (Kirchner, 2004, p. 313). Debuccalization has often been accounted for within an Autosegmental Theory of phonology and a phonological theory that grew out of it: feature geometry. Feature geometry proposes a hierarchal organization of distinctive features instead of traditional matrices, arguing that features can be classified and grouped according to their articulatory function and phonological behavior (i.e., spreading, linking, delinking, etc.), and that each group of features is dominated by a node within the hierarchy. Figure 8, adapted from McCarthy (1988), shows this full geometry with the root node in the top tier of the hierarchal structure. The root node is formed by the features [sonorant] and [consonantal], which are unique because they are the only features that "do not assimilate, reduce, or dissimilate except in conjunction with processes that affect the entire segment" (McCarthy, 1988, p. 97).

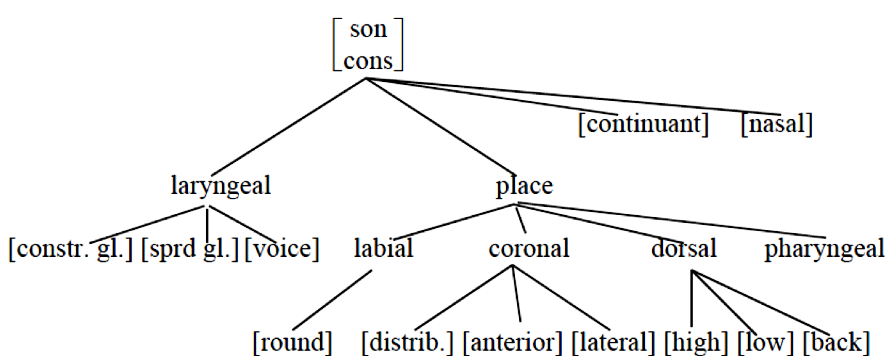

Figure 8. Proposed feature geometry, with each respective node dominating the nodes and/or features on the tier bier below it as indicated, adapted from McCarthy (1988).

Within this theory, we can see that the place features, dominated by the place node, are grouped together because they not only serve a similar articulatory function, but they also behave similarly. For example, these features—treated as autonomous autosegments on a ‘place tier’—can spread to nearby segments, i.e., assimilation. In addition to spreading, these features can disassociate from a segment entirely, a process that involves the delinking of the place node while retaining the features associated with laryngeal specification, i.e., debuccalization. In Spanish, /s/ weakening is easily accounted for within this framework. Figure 9 shows a feature geometrical diagram of the debuccalization process in which the place node of [s] is severed, rendering [h], a segment that is equivalent to [s] in laryngeal specification as well as for the features [continuant] and [nasal].

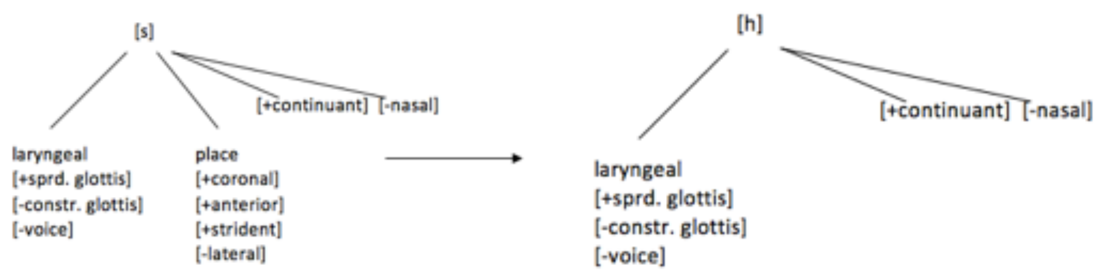

Figure 9. Debuccalization in Spanish: delinking of the place node results in perceived /s/-weakening.

Connecting Figure 9 back to source-filter theory, we can think about the features associated with the laryngeal node as being specified in the source while features in the place node are specified in the filter. When this place node is deleted, the only features that remain as a result of the filter configuration are [+continuant], signifying that air is able to flow through the median plane of the oral tract, and [-nasal], signifying that the velum is raised, blocking airflow through the nose. While source-filter theory is helpful for understanding the phonetic relationship between

[s] and [h], an autosegmental account of debuccalization is similarly helpful in accounting for the phonological relationship between these sounds.

3.2.3 Motivating /s/ lenition

Both phonetically and phonologically, it is clear that [h] is a simplified sound as compared to [s]. While [s] is produced via a complex oral gesture, [h] avoids this gesture entirely, rendering it less physiologically complex and therefore suggesting that less articulatory effort is required for its production (Guitart, 1976, p. 74). The proposal that articulatory economy guides /s/ weakening is not a new one, and when one understands the relationship between [s] and [h] it is clear why this idea has been espoused by so many. However, as Kirchner (2004, p. 313) observes, "previous treatments have failed to offer an empirically adequate, unified formal characterisation of lenition, or to account for the contexts in which lenition typically occurs."

In this dissertation, I propose a phonological analysis of Spanish /s/ weakening inspired by Kirchner's (1998, 2004) effort-based approach in which lenition patterns arise directly from the competing needs to minimize articulatory effort and preserve perceptual distinctions. Crucially, I believe this approach can account for patterns of both onset and coda weakening within a single framework. Before discussing Kirchner's proposal and how it might be adapted and subsequently applied to sibilant weakening in Salvadoran Spanish, the following sections provide an overview of two theories on which Kirchner's approach is predicated: Optimality Theory and phonetically-based phonology.

3.3 Optimality Theory

Optimality Theory (Prince & Smolensky, 1993/2004) posits that surface forms of language reflect the resolution of conflicts between competing demands. Unlike derivational phonology, a framework in which a series of rules formulated without clear motivation are applied to an underlying representation to generate an observed surface form, OT sees language as a system of conflicting forces (called *constraints)* with typological or phonetic motivations. Within OT, constraints are considered to be:

Universal: Constraints are not specific to a given language or group of languages. Every language and dialect functions under the same set of universal constraints.

Ranked: Cross-linguistic (or cross-dialectical) differences are accounted for by differences in how universal constraints are ranked.

Violable: All surface forms, called *outputs*, violate some constraints.

In conflict: 'Optimal' outputs are those that are most 'harmonic' within a given constraint ranking. That is, they have the least serious violations of the ranked set of constraints for a language and/or dialect.

OT-style constraints are divided into two sub-types: *Markedness* and *Faithfulness*. Markedness constraints are concerned with the well-formedness of the output and are therefore in conflict with the input (i.e., the underlying form) if it contains a 'marked,' or cross-linguistically unfavorable, form of some kind. Markedness constraints tend to reflect cross-linguistic typological patterns such as a dispreference for coda consonants, which is formalized within OT as NOCODA ('syllables should be open'). Faithfulness constraints, then, are in direct conflict with their Markedness counterparts in that they enforce similarity between input and output forms with the goal of maintaining continuity and limiting how much change can occur. For example, in direct conflict with the NoCoda Markedness constraint might be IO-Max-C, specifying that consonants should not be deleted between the input and output (a strategy that could solve our coda problem). In this way, Markedness and Faithfulness constraints work as a system of checks and balances: a language should avoid marked structures whenever possible without obscuring necessary similarities between the input and output.

Crucially, two forces facilitate the conversion of input forms into their respective output forms: Gen and Eval. First, Gen generates a potentially infinite list of output candidates for a given input. Next, Eval applies the pertinent constraint ranking to the list of possible outputs and selects the most harmonic candidate. This process is illustrated in Figure 10 below.

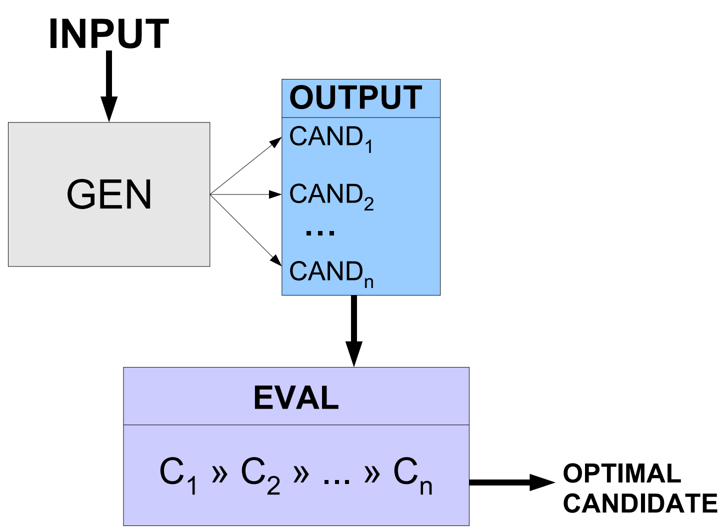

Figure 10. OT schematic from Schulz (2007).

In a standard OT analysis, a list of the most relevant candidates is displayed vertically in a table called a *tableau,* with relevant constraints listed horizontally along the top. Constraint violations incurred by candidates are denoted by an asterisk '*', winners (i.e., optimal candidates) are marked with a pointed finger or arrow, '→', and an exclamation point '!' indicates that a given candidate has been eliminated and by which constraint. Shaded cells indicate irrelevance once a candidate has been selected by EVAL. Table 1 below shows a simple tableau in which the Markedness constraint from the previous example, NOCODA, interacts with the Faithfulness constraint IO-MAX-C with respect to Spanish /s/ lenition in the word *mosca* /moska/ ('fly'). In this /s/ weakening dialect, NOCODA is ranked above IO-MAX-C, so the winning candidate is the one in which the coda /s/ from the input is deleted.

Table 1

Sample OT tableau for a dialect of Spanish in which NOCODA >> IO-MAX-C

/moska/	NOCODA	IO-MAX-C
a. [mos.ka]	*!	
b. → [mo.ka]		*

Here, we see that EVAL selects the unfaithful [mo.ka] as the optimal candidate because avoiding coda consonants is more important than maintaining consonants between the input and the output. Crucially, while the winner [mo.ka] violates the Faithfulness constraint IO-MAX-C, the violation occurs after an optimal candidate has already been selected, as indicated by shading.

Contrastingly, in another dialect of Spanish in which /s/ is not lenited, the constraint ranking would be reversed and the faithful [mos.ka] would be selected as the winning candidate, as in Table 2 below. Again, while [mos.ka] violates the Markedness constraint NOCODA, it has already been selected as the winning candidate by the time this violation occurs.

Table 2

Sample OT tableau for a dialect of Spanish in which IO-MAX-C >> NOCODA

/moska/	IO-MAX-C	NOCODA
a. → [mos.ka]		*
b. [mo.ka]	*!	

In this sense, OT can be a useful framework for thinking about dialectal differences in the Spanish-speaking world which, as detailed in previous sections, is typically divided into phonologically innovative and phonologically conservative dialects. While these dialects have the same fundamental constraints at work, a preference for Faithfulness over avoiding Markedness (exemplified by *las tierras altas*) or vice versa (exemplified by *las tierras bajas*) speaks to the major differences observed between the two typologies.

However, as we saw in the previous chapter, patterns of Spanish /s/ lenition are conditioned by a myriad of factors and reducing them to generalizations is a dangerous oversimplification. That is, even in the most conservative, /s/-retaining dialects, speakers do not categorically produce the sibilant in all linguistic contexts, nor do speakers in even the most /s/-reducing dialects weaken the sibilant at all cost. This language-internal variation is further complicated by factors affecting inter- and intra-speaker variation, such as those outlined in the previous chapter. These nuances are, admittedly, difficult to capture within an OT framework; this issue, which is one of the chief concerns of this dissertation, will be revisited in Section 3.6.

The proposal that grammars are built on a set of ranked constraints governed by Markedness laws has sparked a fascinating conversation about where these Markedness laws originate. The following section examines the hypothesis that a speaker's knowledge of Markedness comes from his or her phonetic knowledge, a theory that lays the foundation for Kirchner's (2004) effort-based approach to consonant lenition.

3.4 Phonetically-based phonology

In an effort to explore the question of where and how Markedness constraints originate, 2004's *Phonetically Based Phonology* (Hayes et al., 2004) hypothesizes that "phonological constraints can be rooted in phonetic knowledge (Kingston & Diehl, 1994), and speakers' partial understanding of the physical conditions under which speech is produced and perceived" (Hayes & Steriade, 2004, p.1). This theory about the interface of phonological Markedness and phonetic knowledge was first posited in Chomsky and Halle's *The sound pattern of English* (1968); however, within an OT framework, it supposes that constraints embody some kind of phonetic difficulty, either in production or perception. While all speaker/hearers will experience similar phonetic difficulties—causing them to posit a set of universal constraints resulting in phonological typologies—not all languages will remedy these issues in the same way. Constraints that are induced from evidence about phonetic conditions are referred to as *markedness from phonetics*, and sound patterns reflect the activity of these constraints.

A common example of Markedness from phonetics can be seen in the cross-linguistic tendency to avoid voiced obstruent geminates, embodied in the following constraint:

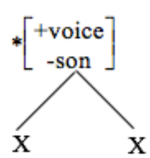

Hayes and Steriade (2004, p. 8) argue that this pattern can be attributed to the aerodynamics of voicing: "active oral tract expansion (for example, by tongue root advancement or larynx lowering) is necessary to maintain airflow by an obstruent. These maneuvers cannot be continued indefinitely or controlled tightly. It is therefore more difficult to sustain production of voicing in long obstruents." Similarly, the authors cite aerodynamic restrictions to explain asymmetries in singleton voiced obstruents observed by Ohala and Riordan (1979), who find that voicing in bilabial [b] is longer (82 ms) than in coronal [d] (63 ms) and velar [g] (52 ms): "the larger cavity behind the lips offers more compliant tissue, which allows the cavity to continue for a longer time to expand passively in response to airflow" (Hayes & Steriade, 2004, p. 8). This voicing example illustrates how Markedness constraints arise as a direct result of speaker/hearer knowledge of phonetic difficulty.

Phonetic difficulty posed by aerodynamics can also serve to explain, at least in part, the sample Markedness constraint from the previous section: NOCODA. In the example presented in Section 3.3, Spanish /s/ weakening was driven by a dispreference for consonants in the syllable coda, a position that poses phonetic challenges compared to the syllable onset. For example, in an experimental investigation of the differences in aerodynamic characteristics for onset and coda fricatives, Solé (2003) finds that coda fricatives have a slower build-up of oral pressure and a lower pressure peak as compared to onsets, making it more difficult for them to generate audible turbulence. She notes that fricatives in general "allow lesser articulatory and aerodynamic variation than other segment types, and they are more likely to decay if the precise postural and aerodynamic requirements are not met" (Solé, 2003, p. 2761). Again, in this example we see a clear phonetic motive for a common Markedness constraint.

Beyond this obstruent voicing example, which the authors examine in great detail, Hayes and Steriade (2004) review additional markedness scales with phonetic bases, such as *scales of perceptibility*. Jun (2004) utilizes this scale in an account in which he argues that place assimilation is a result of conflict between effort avoidance (i.e., the desire to eliminate or reduce any articulatory gesture) and the preservation of forms that are perceptually sensitive. Within this account, "less perceptible gestures are more likely to disappear" (Hayes & Steriade, 2004, p. 19).

While the idea that phonetics and phonology are inextricably linked is undisputed, some scholars take a different stance on the role of phonetics in phonology that is in conflict with that proposed by Hayes et al. Blevins (2004), for example, argues for a theory of evolutionary phonology in which diachronic patterns play the most important role in determining phonological typology. Crucially, from this perspective, learners do not impose what they know about phonetics on phonological structure; the role of phonetics is simply to determine which sound changes are acceptable and frequent. In their chapter that appears in the Hayes and Steriade volume, Blevins and Garrett (2004) argue that the principal cause of metathesis-related sound change is listener-based reinterpretation, known as "innocent misapprehension" (Ohala, 1981, 1990). Innocent misapprehension encompasses the idea that learners mishear forms and consequently assign new structural interpretations to them, initiating sound change. New forms that persist diachronically are phonetically effective not because learners are using their phonetic knowledge to create grammars, but because only phonetically effective phonologies are able to survive.

The most controversial aspect of this view, Hayes and Steriade note, is not that diachrony plays some role in explaining phonological naturalness, but that a diachronic account alone can account for phonological patterns in the world's languages. Hayes and Steriade's argument, which I adopt in this dissertation, is that a theory of phonology cannot be entirely blind to phonetically-based Markedness principles. To justify their position, the authors refer to empirical studies whose results come in direct conflict with Blevins's view. For example, they note that innocent misapprehensions are unable to account for studies of child phonology in which children make changes to adult language forms that they hear accurately (e.g., Smith, 1973). They also express doubt that innocent misapprehension is directionally stable enough to drive diachronic sound change, citing studies of perceptual assimilation in which nasals are often misperceived non-assimilatorily (e.g., Hura, Lindblom, & Diehl, 1992), a finding that contradicts what evolutionary phonology would predict.

Assuming the hypothesis that phonologies are determined by phonetic principles, the following section describes Kirchner's (2004) effort-based account of consonant lenition, which forms the basis of the analysis presented in this dissertation, and examines its utility in accounting for Salvadoran /s/ lenition.

3.5 An effort-based approach to consonant lenition

Consonant lenition manifests in a number of different ways. Observed both in synchronic alternations as well as diachronic sound change, consonantal lenition processes are prevalent cross-linguistically; such processes include degemination (reduction of a long to a short consonant, e.g. tː → t), flapping (the reduction of a stop to a flap, e.g. t → ɾ), spirantization (the reduction of a stop or affricate to a fricative or approximant continuant, e.g. d͡ʒ → ʒ), debuccalization (reduction to a laryngeal consonant, as in s → h), and voicing (e.g. t → d). Kirchner (2004, p. 313) writes that a "unified characterisation of... 'weakening' has been a vexed question of phonological theory; but the core idea, as applied to consonants, is some *reduction of constriction degree or duration*" [emphasis original]. What unifies these processes, then, is that a) 'weak' is used a comparative term in which a given sound is weaker than the sound from which it came, and b) the resulting (weaker) sound is less articulatorily complex, and therefore easier to produce, than the original (Kirchner, 2004; Zuraw, 2009).

With the goal of proposing a framework that can account for all instances of lenition as well as the contexts in which they occur, Kirchner (1998, 2004) offers an articulatory account in which lenition arises directly from the goal of minimizing articulatory effort. The basic idea is this: the more constricted a consonant, the greater effort needed to produce it; a consonant with more constriction must move the articulator a greater distance to reach its target and requires greater velocity (and therefore energy) to move the articulator that distance in a given amount of time. In an OT account, EVAL must compute the total *effort cost*, which Kirchner (2004, p. 314) operationalizes as "a mental estimate of the biomechanical energy required for articulatory production of the candidate," for each output candidate generated by GEN. Candidates with higher effort cost incur more violations for phonetically-driven Markedness constraints, termed LAZY. Consistent with the phonetically-based phonological framework, Markedness derived from the imperative to minimize articulatory effort cost interacts with a need to preserve perceptual distinctions. The sections that follow describe relevant aspects of Kirchner's approach and how they will be applied to my own analysis.

3.5.1 Effort cost as Markedness

Kirchner develops a series of Markedness constraints in the form LAZY (C, K, R), where C is a class of consonants, K is the context in which they occur, and R is rate of speech. This proposal is based on the idea that lenition is gradient, and that the degree to which a consonant is lenited depends on consonant class, phonological context, and speech rate.[6]

Based on inferences about the relative effort cost of different consonant types, which follow from equations of effort based on biomechanical energy (Lindblom, 1983), Kirchner first develops a hierarchy of LAZY constraints. This hierarchy reflects the assumption that, the more effortful the gesture, the more impetus to lenite. He connects this conception of "effortfulness" directly to both degree of constriction achieved and time spent with the articulators in their most displaced position, i.e., sustained partial constriction. Kirchner illustrates this idea by comparing stops, non-strident, and strident fricatives. Of these three consonant classes, strident fricatives are the most effortful of all, as they "involve a sustained interval of precise, close constriction" which "requires antagonistic muscle activation, in opposition to the (agonist) closure activation, to bring and keep the articulator in close constriction, without allowing it to go all the way to closure." This is in contrast to non-strident fricatives, which are produced in such a way that the active articulator can approach the passive articulator with little purpose or precision. Figure 11 and Figure 12, taken from Kirchner (2004, p. 327), illustrate this comparison visually.

[6] For reasons of feasibility, this dissertation will not examine the effect of speech rate on Salvadoran /s/ lenition. Therefore, this factor is not reviewed here.

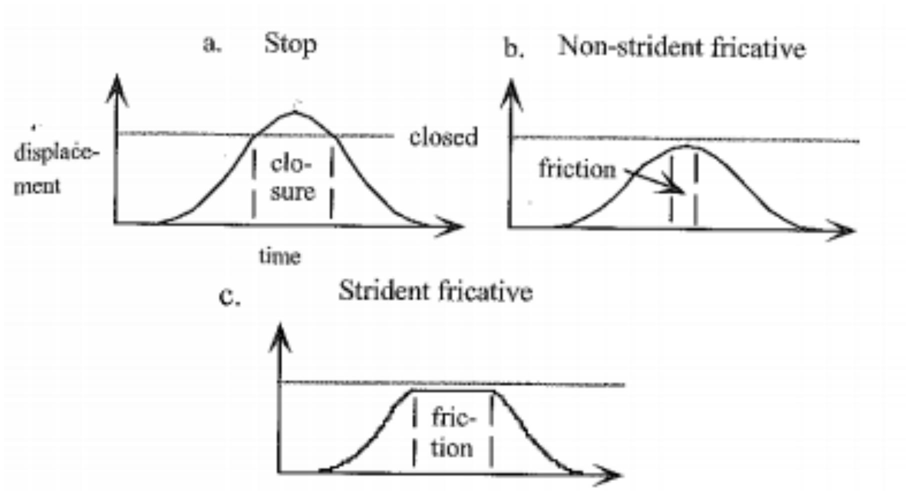

Figure 11. Kirchner (2004, p. 327) compares constriction in stops, non-strident, and strident fricatives, showing that strident fricatives require sustained, close constriction over an interval of time.

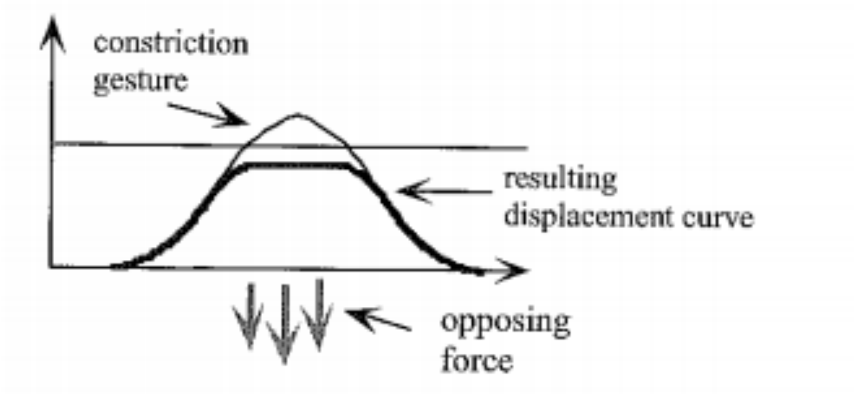

Figure 12. Kirchner (2004, p. 327) schematizes the agonist forces at work in the articulation of strident consonants, rendering them the most effortful of the consonant classes compared here.

The assertion that strident fricatives are more effortful than non-strident fricatives can, crucially, be captured by the following ranking of Kirchner's effort-based Markedness constraints: LAZY (vcl/vcd_strid_fric) >> LAZY (vcl/vcd_fric).

In addition to consonant class, the openness of the segments flanking a given consonant affect its effort cost within Kirchner's framework. This is logical, as "the more open the flanking segments, the greater the displacement (hence effort) required to achieve a given degree of constriction” (Kirchner, 2004, p. 315). This relationship between displacement, degree of constriction, and flanking segments can be visualized in Figure 13 below.

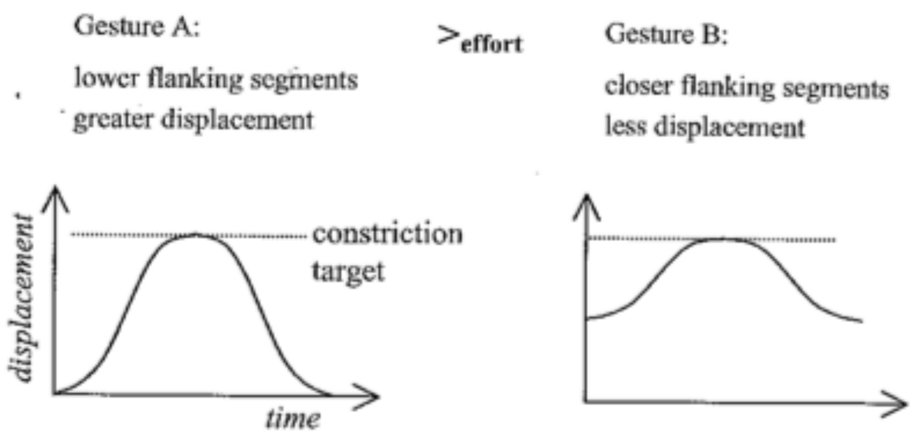

Figure 13. Kirchner (2004, p. 323) schematizes the relationship between displacement, degree of constriction, and flanking segments, showing that more open flanking segments require more displacement to reach the specified degree of constriction and are therefore more effortful.

Following this schematic, Kirchner proposes that intervocalic position is more effortful than non-intervocalic position in which there is only one jaw displacement gesture. This can be formalized with respect to markedness as: LAZY (C, V_V) >> LAZY (C, C_V), LAZY (C, V_C). Furthermore, jaw displacement in consonant constriction relative to the flanking vowel(s) should also play a role, with open vowels such as /a/ requiring more displacement than closed vowels like /i/ and therefore encouraging more lenition. Kirchner formalizes this gradient situation in his analysis of Florentine Italian by distinguishing between flanking high vowels and flanking non-high vowels, such that: LAZY (C, $V_{[-high]}$_$V_{[-high]}$) >> LAZY (C, $V_{[+high]}$_$V_{[-high]}$), LAZY (C, $V_{[-high]}$_$V_{[+high]}$) >> LAZY (C, $V_{[+high]}$_$V_{[+high]}$) and LAZY (C, C_$V_{[-high]}$), LAZY (C, $V_{[-high]}$ _C) >> LAZY (C, C_$V_{[+high]}$), LAZY (C, $V_{[+high]}$ _C).

In his discussion of the factors that affect biomechanical articulatory effort, Kirchner draws a number of comparisons, the most relevant of which have been discussed in this section. Table 3 summarizes these comparisons.

Table 3

Effort cost comparisons by articulatory factor

Articulatory factor	More effortful	Less effortful
Vertical displacement of the articulators active in consonant constriction	Greater displacement, as in a strident fricative	Lesser displacement, as in a stop
Number of jaw displacement gestures	Two gestures: C-to-V and V-to-C, as in VCV	One gesture: C-to-V or V-to-C, as in (C)CV, VC(C)
Jaw displacement in consonant constriction relative to flanking vowel	Greater displacement, as in $V_{[-high]}$ //strident fricative	Lesser displacement, as in $V_{[+high]}$ //strident fricative

Factors affecting effort cost will be revisited in the context of Salvadoran /s/ lenition later, in Section 3.6.

3.5.2 Preservation of perceptual distinctions

Interacting with phonetically-based markedness in Kirchner's account of effort-based lenition is Faithfulness in the form of lenition-blocking constraints. This is one of the most crucial aspects of Kirchner's account, which is able to account for all types of lenition because the type of change that occurs depends on which lenition-blocking constraints interact with LAZY. For example, in the case of spirantization, LAZY interacts with a faithfulness constraint that requires the preservation of an input [-continuant] specification, as seen in Table 4a and 4b.

Table 4

Sample OT tableaux from Kirchner (2004, p. 314) in which spirantization is either allowed (a) or blocked (b) depending on the ranking of LAZY and PRESERVE(continuant) ('preserve the input [-continuant] specification in the output').

(a)

/d/	LAZY	PRES(cont)
a. [d]	**!	
b. → [ð]	*	*

(b)

/d/	PRES(cont)	LAZY
a. → [d]		**
b. [ð]	*!	*

In Table 4a, reducing effort cost is prioritized even at the expense of the perceptual distinction between consonants that are [-continuant] and [+continuant]. Contrastingly, in Table 4b, preservation of this particular perceptual distinction is more important than modulating effort cost, and spirantization via reduction of the articulatory gesture is blocked. It should be noted that, in these sample tableaux, output candidate [d] incurs two violations for LAZY while candidate [ð] incurs only one because stops are considered to be more effortful than continuants given the distance traveled by the articulator (compare Figure 12 and Figure 13).

It should be noted that the input [d] and output candidate [ð] in this example only differ with respect to one feature, [continuant], rendering a fairly straight-forward formulation of the Faithfulness constraint. Kirchner confronts more complex cases, such as accounting for the weakening of the strident /s/ to the non-strident approximate [s̹] without full debuccalization to [h]. This poses a challenge because both [s̹] and [h] violate the Faithfulness constraint PRESERVE(strident) ('preserve the input [+strident] specification in the output'). In order to differentiate these two outcomes, Kirchner incorporates what he calls a 'disjunctively combined

Faithfulness constraint' that penalizes candidates with a different specification for both [strident] and [place].

Furthermore, Kirchner adds that, just as faster speech rates increase the effort cost of a given gesture, lower (i.e., less formal) speech registers likewise give rise to more lenition as the speaker shifts toward hypoarticulation. He formalizes this as LENITION-BLOCKING-CONSTRAINT/REG >> LENITION-BLOCKING-CONSTRAINT/REG', where *Reg* represents a more formal register and *Reg'* represents a less formal register.

The following section describes now Kirchner's framework will be applied to my analysis of Salvadoran /s/ lenition.

3.6 An effort-based approach to Salvadoran /s/ lenition

Kirchner's (2004, p. 316) argument that an effort-based approach can account for the fact that "lenition indeed constitutes a unified phenomenon, not an arbitrary collection of unrelated processes" makes his framework ideal for evaluating Salvadoran /s/ lenition, a process that is arguably more varied and complex than in other dialects of Spanish. In addition to debuccalization and deletion of /s/ in coda position, Salvadoran Spanish shows pervasive evidence of the following:

(a) Debuccalization and deletion of /s/ in syllable-onset position, both word-medially and word-initially;

(b) The weakening of /s/ to a non-strident approximant perceptually similar to [θ][7]; and

[7] The treatment of this allophone as the product of weakening is likely a controversial one, as it has traditionally been treated as a historical artifact in Central American dialects of Spanish. Chapter 4 is, therefore, devoted to a discussion of this sound, including acoustic and phonological analyses presented as evidence for lenition.

(c) Variable voicing (sometimes accompanied by debuccalization) of /s/.[8]

It appears that there are three disparate lenition processes affecting Salvadoran /s/, and I propose that an effort-based approach within a phonetically-based phonological perspective can account for all of them in one, unifying analysis.

It is important to note that, while voicing of Spanish /s/ is not always treated as weakening in the literature, obstruent voicing is another manifestation of lenition. Kirchner (2004, p. 313) notes that, "although ostensibly involving an adjustment in laryngeal specification rather than reduction of constriction, [voicing] is also standardly included in this [lenition] typology." Kirchner (2004, p. 326) explains the effort cost associated with obstruent voicing via aerodynamic effects proposed by Westbury and Keating (1986), who show that singleton stops preceding a voiced sonorant in utterance-medial position "undergo *passive voicing*, unless they are devoiced by active abduction (or constriction) of the glottis, assuming an adducted rest position of the glottis." Kirchner adds that his own simulations show that the same is true for medial singleton continuant consonants, resulting in the following ranking if all else is equal: LAZY (vcl_strid_fric, V_V) >> LAZY (vcd_strid_fric, V_V). Within this framework, [s] is more effortful than [z]. Consistent with this assertion, accounts of variable /s/ voicing in Spanish find that its use is conditioned by factors that also condition debuccalization and deletion, such as prosodic position (Davidson, 2014; McKinnon, 2012), speech rate, and syllable stress (Davidson, 2014; García, 2015).

[8] Intervocalic voicing of /s/ has been documented in a number of Spanish dialects, including those of Spain (Davidson, 2014; McKinnon, 2012; Torreblanca, 1986), Ecuador (Aguirre, 2000; Calle, 2010; García, 2015; Lipski, 1989), and Costa Rica (Chappell, 2016).

3.6.1 Effort-based Markedness constraints

As we saw in saw in Figure 12 above, strident consonants like /s/ require precise and prolonged constriction via agonist force in order to achieve their characteristic turbulence. Kirchner (2004, p. 326) explains: "In sibilants in particular, it is known that such antagonism is required: namely, the tongue-blade constriction is partially opposed by a stiffening and bracing of the sides of the tongue against the molar gumline, to produce a grooved channel for the airflow." Given the difficulties associated with producing /s/, it is hardly surprising that Salvadoran Spanish employs so many strategies to mitigate its effort cost. These strategies include voicing (/s/ → [z]), gestural undershoot (/s/ → [s^{θ}]), debuccalization (/s/ → [h]), debuccalization plus voicing (/s/ → [ɦ]) and deletion (/s/ → [Ø]).

In this section, I present the Markedness constraints that will be employed in my analysis of Salvadoran /s/ lenition. These constraints are modeled after those proposed by Kirchner but modified in accordance with the goals of this dissertation. While Kirchner's constraints take the form of LAZY(C, K, R), where *C* is a class of consonants, *K* is the context in which /s/ occurs, and *R* is speech rate, Markedness constraints in this dissertation are style as LAZY(G, K), where *G* is a gesture or combination of gestures and *K* is context. Constraint types are defined (1)-(5) below:

(1) LAZY(vcl_strid_fric, K): This constraint penalizes the simultaneous production of glottal abduction and the antagonistic forces required at the place of articulation for strident consonants in context K.

(2) LAZY(strid_fric, K): This constraint penalizes the production of the antagonistic forces required at the place of articulation for strident consonants in context K.

(3) LAZY(oral_gest, K): This constraint penalizes the production of any oral gesture in context K.

(4) LAZY(glottal_abd, K): This constraint penalizes the production of glottal abduction in context K.

(5) LAZY(any_gest): This constraint penalizes the production of any gesture and is not constrained by context.[9]

Calculations of biomechanical effort cost give rise to following ranking conditions for these five constraint types: LAZY(vcl_strid_fric, K) >> LAZY(strid_fric, K) >> LAZY(oral_gest, K) >> LAZY(any_gest); LAZY(vcl_strid_fric, K) >> LAZY(glottal_abd, K) >> LAZY(any_gest). In other words, if some oral gesture is produced, LAZY(any_gest) and LAZY(oral_gest, K) are violated; if that oral gesture is a strident fricative, LAZY(strid_fric, K) is additionally violated; and if that gesture is accompanied by glottal abduction, LAZY(vcl_strid_fric, K) is also additionally violated. Similarly, if glottal abduction is produced, both LAZY(any_gest) and LAZY(glottal_abd, K) are violated; if that glottal gesture is accompanied by an oral gesture, LAZY(vcl_strid_fric, K) is additionally violated. Table 5 shows the implications of these ranking conditions for the six variants of interest.

[9] The premise of reducing effort cost in specific contexts is predicated on the assumption that certain gestures are more difficult to execute in those contexts. If the constraint does not target a specific gesture, there are no more or less difficult contexts.

Table 5

Markedness constraints and the processes they affect

Process	LAZY (vcl_strid_fric)	LAZY (strid_fric)	LAZY (oral_gest)	LAZY (glottal_abd)	LAZY (any_gest)
/s/ → [s]	*	*	*	*	*
/s/ → [z]		*	*		*
/s/ → [s^{θ}]			*	*	*
/s/ → [h]				*	*
/s/ → [ɦ]					*
/s/ → [∅]					

Crucially, [z] and [s^{θ}] violate the same number of Markedness constraints as I am not making any assertions about the difficulty of glottal abduction versus the difficulty of the gesture required to produce a strident fricative.

The second component of effort-based Markedness constraints, context, is much more complex, particularly because I am interested in the behavior of /s/ in all prosodic positions. Furthermore, it is important to note that the difficult of a given context differs depending on the gesture in question, so I will first discuss contexts affecting LAZY(vcl_strid_fric), LAZY(strid_fric), and LAZY(oral_gest). According to an effort-based framework, LAZY(vcl_strid_fric), LAZY(strid_fric) should be equally impacted by context with respect to the oral gesture, while LAZY(oral_gest) should be impacted slightly less. I begin with possible intervocalic contexts which, according to Kirchner's typology, are more poised for lenition than non-intervocalic contexts due to the openness of the flanking segments. Constraints according to context are listed below in (6)-(8), ordered from most effortful to least effortful:

(6) LAZY(G, $V_{[-high]}$__ $V_{[-high]}$)

(7) LAZY(G, $V_{[-high]}$___ $V_{[+high]}$), LAZY(C, $V_{[+high]}$___ $V_{[-high]}$)

(8) LAZY(G, $V_{[+high]}$___ $V_{[+high]}$)

This ranking reflects the ordering of possible intervocalic contexts in which /s/ occurs in terms of effort cost. Because Kirchner makes no assertion as to whether the height of the preceding or following vowel has more of an effect on the effort required to articulate the intervening sibilant, the relationship between these constraints as seen in (7) cannot be determined at this time. Again, these contexts should pose more difficulty for LAZY(vcl_strid_fric) and LAZY(strid_fric) than for LAZY(oral_gest), as the latter does not require the tongue to come into close or sustained contact with the articulatory target.

With respect to C__V contexts, it is important to note that Kirchner does not deal with these explicitly beyond the assertion that the number of jaw displacement gestures is less in these contexts, rendering them less effortful. When /s/ occurs as a syllable-onset following another consonant, as in the word *bolsa* /bolsa/ ('bag'), the more constricted preceding segment reduces the displacement required by the active articulator to achieve the sibilant. Therefore, constraints for C__V contexts are again ranked according to vowel height, as can be seen in (9)-(10):

(9) LAZY(G, C__$V_{[-high]}$)

(10) LAZY(G, C__$V_{[+high]}$)

Again, these contexts should pose more difficulty for LAZY(vcl_strid_fric) and LAZY(strid_fric) than for LAZY(oral_gest) and should have no effect on LAZY(glottal_abd).

While Kirchner (2004) does not explicitly address preconsonantal contexts, aerodynamic differences between onset and coda fricatives result in a number of asymmetries such as a slower build-up of oral pressure and a lower pressure peak (see Section 3.4). In addition to the inherent effortfulness of [s], this reduction in pressure poses cumulative challenges to producing audible frication. Furthermore, it is likely that the class of the consonant that follows /s/ has an effect on lenition rates, as Kirchner (2004, pp. 324–325) notes that "any factor that would raise the effort cost in a given context boosts the probability of lenition occurring in that context" and that "if we instead adopt a more biomechanically plausible gradient set of contextual distinctions in effort cost, the accuracy of the analysis would actually improve." With the effects of displacement in mind, it is reasonable to assume that following consonants that share gestures with [s] are likely to encourage retention of the voiceless sibilant, as less effort is required in the CC transition. Specifically, consonants that share specifications for place (i.e., [+coronal]) might facilitate this transition by reducing the distance the articulator needs to travel to achieve the required degree of constriction. Such a prediction finds support with Jun (2004), who argues that place assimilation is the result of effort avoidance. Similarly, if /s/ occurs before a voiced consonant, two transitions are required by the glottis: it must actively abduct between the preceding vowel and [s], and then return to its adducted rest position for the following consonant. However, if the following consonant is voiceless, the glottis is only required to make the single transition, from the vowel to the voiceless sibilant. This prediction is corroborated by Campos-Astorkiza (2014), who argues that sibilant voicing assimilation in Spanish is the result of gestural blending, which is a manifestation of weakening.

Taking all of these factors into account, I propose the constraints in (11)-(19) for preconsonantal contexts. This set of contexts presents more complexities than the previous two, as all four constraint types should be affected slightly differently. First, I hypothesize that LAZY(vcl_strid_fric), LAZY(strid_fric), and LAZY(oral_gest) will be modulated by both preceding vowel height and coronality of the following consonant. Second, I hypothesize that LAZY(vcl_strid_fric) and LAZY(strid_fric) will differ according to voicing, with the importance of the latter decreasing before voiced consonants because of voicing agreement with [z]. Finally, I hypothesize that LAZY(glottal_abd) will be conditioned only by voicing of the following consonant, with elevated importance before voiced consonants. These contexts are listed below, in no particular order:

(11) Lazy(G, $V_{[-high]}\_C_{[-cor, +voice]}$)

(12) Lazy(G, $V_{[+high]}\_C_{[-cor, +voice]}$)

(13) Lazy(G, $V_{[-high]}\_C_{[-cor, -voice]}$)

(14) Lazy(G, $V_{[+high]}\ \_C_{[-cor, -voice]}$)

(15) Lazy(G, $V_{[-high]}\ \_C_{[+cor, +voice]}$)

(16) Lazy(G, $V_{[+high]}\ \_C_{[+cor, +voice]}$)

(17) Lazy(G, $V_{[-high]}\ \_C_{[+cor, -voice]}$)

(18) Lazy(G, $V_{[+high]}\ \_C_{[+cor, -voice]}$)

(19) Lazy(G, C _C)[10]

Thus far, this section has introduced 15 surrounding phonological contexts that are hypothesized to modulate patterns of effort-based /s/ lenition. The final set of contexts to examine are those in which /s/ is adjacent to a pause, occurring either phrase-initially or phrase-finally. As we can extrapolate from Table 3, these contexts involve a similar amount of articulatory displacement to C__V contexts, as the principle jaw displacement gesture occurs between the preceding or following vowel and /s/. Again, I hypothesize that the importance of constraints LAZY(vcl_strid_fric), LAZY(strid_fric), and LAZY(oral_gest) in pause-adjacent contexts will be highly dependent on the height of the adjacent vowel, with LAZY(vcl_strid_fric) and LAZY(strid_fric) outranking LAZY(oral_gest). This last set of contexts, listed in order from easiest to most difficult with respect to these three constraint types, appears in (20)-(21):

[10] For this hybrid context, which is both post- and preconsonantal, I have not made any assumptions other than that this context will be easier than other preconsonantal contexts given the degree of constriction specified for both preceding and following segments. The only possible CCC sequences in Spanish are /ksp/, /kst/, /nsp/ and /nst/, and tokens of this nature occurred very infrequently in my Salvadoran data (*n* = 18, or <1% of tokens).

(20) Lazy(G, $V_{[-high]}$__##), Lazy(G, ##__$V_{[-high]}$)

(21) Lazy(G, $V_{[+high]}$__##), Lazy(G, ##__$V_{[+high]}$)

In sum, this section has proposed five constraint types based on gesture, four of which are further conditioned by phonological context. Additionally, I have hypothesized that the nature of the gesture penalized by each constraint type will determine the difficulty posed by a given context. Now that Markedness constraints have been established, the following section introduces rival Faithfulness constraints in the grammar.

3.6.2 Perceptually-based Faithfulness constraints

While Kirchner formulates his faithfulness constraints as LENITION-BLOCKING-CONSTRAINT/REG, I will not take register into account as this dissertation is exclusively interested in casual, naturalistic speech. Instead, to account for linguistic factors that have a proven effect on perceptual salience, I incorporate prosodic position and syllable tonicity. Taking these factors into account, the Faithfulness constraints proposed for my analysis of Salvadoran Spanish are formulated as follows: LENITION-BLOCKING-CONSTRAINT/ PP, T, where *PP* is prosodic position, and *T* is tonicity. Because lexical frequency should also play a role in the need to preserve perceptual distinctions (see Section 2.3.2.4), I also evaluate its role in Salvadoran /s/ lenition in Section 6.2 in order to determine whether or not this factor should be incorporated into Faithfulness constraints, thereby modifying the form to LENITION-BLOCKING-CONSTRAINT/ PP, T, F, where *F* is a measure of lexical frequency.

This section describes the formulation of these perceptually-based Faithfulness constraints factor-by-factor. First, lenition-blocking-constraints were developed based on the conception of

/s/ as a segment that is [-voice], [+strident], and [+coronal]. Relevant faithfulness constraint types, categorized by feature, appear in (22)-(26):

(22) PRESERVE(strident, voice): Assign one violation for every output form that does not preserve both input [+strident] and [-voice] specifications.

(23) PRESERVE(strident): Assign one violation for every output form that does not preserve the input [+strident] specification.

(24) PRESERVE(coronal): Assign one violation for every output form that does not preserve the input [+coronal] specification.

(25) PRESERVE(voice): Assign one violation for every output form that does not preserve the input [-voice] specification.

(26) PRESERVE(segmental): Assign one violation for every output form that does not preserve the input [+segmental] specification.

The processes that incur violations for each of these constraints are illustrated in Table 6 below.

Table 6

Lenition-blocking Faithfulness constraints and the processes they affect

Process	PRES (strid, voice)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)
/s/ → [s]					
/s/ → [z]	*			*	
/s/ → [s^{θ}]	*	*			
/s/ → [h]	*	*	*		
/s/ → [ɦ]	*	*	*	*	
/s/ → [Ø]	*	*	*	*	*

As the reader can see, the more constraints violated by an output, the higher its perceptual cost. While [z] and [s^{θ}] are the least-costly of the unfaithful outputs, each involving only one substantial feature change, [∅] is the costliest with respect to perceptual salience, violating all five Faithfulness constraints. Again, with respect to [z] and [s^{θ}], I'm making no comparative assertions here about the perceptual cost of changing the input specification for [+strident] and [-voice]. In the eyes of the grammar, both outputs are equally (un)faithful.

The second factor to consider in the formulation of Faithfulness constraints is prosodic position. To understand the effects of prosodic position on consonant weakening, it is crucial to understand the relationship between prosody and the phonetic encoding of segments. The basic idea is this: the phonetic realization of an individual speech segment depends, in part, on its position within the larger prosodic structure of the intonational phrase in which it occurs. Crucial to the present analysis is the idea that the strength of a given prosodic position will have an effect on the strength of the segment that occurs there (see Section 2.3.2.2). By taking into account the prosodic position in which a given /s/ occurs, then, we can identify that context as prosodically strong (i.e., prominent) or weak relative to other contexts.

In the present analysis, I adopt Keating's (2006, p. 171) proposal of domain-initial effects: prosodic strengthening is cumulative, so "the higher in the prosodic tree an initial position is, the stronger that position and the segment in it." Therefore, with respect to syllable-onset /s/, prosodic strength can be assessed using the following scale: phrase-initial > word-initial > syllable-initial. It should be noted that here and throughout this dissertation, *prosodic position* refers to "the highest domain which some segment is initial" (Keating, 2006, p. 171), so *word-initial* denotes any /s/ that is word-initial but not word-final, and *syllable-initial* denotes any /s/ that is syllable-initial but not word-initial. With respect to syllable-final /s/, prosodic strength will be assessed according to initial-ness within the word, with word-medial coda position considered prosodically stronger than word-final coda position. This gives way to the following hierarchy: phrase-initial > word-initial

> syllable-initial > syllable-final > word-final.

Due to the well-attested phonetic effect of phrase-final lengthening, which serves to mark important prosodic boundaries for the listener, phrase-final position does not fall neatly into the hierarchy outlined above. Furthermore, previous studies on Spanish /s/ weakening have yielded inconsistent results in this position, with some reporting high rates of deletion and others reporting high rates of retention. For this reason, I do not make any hypotheses about phrase-final position within this framework beyond it dominating word-final (phrase-medial) position with respect to prosodic strength.

Also related to prosody and a key element of perceptual prominence is syllable tonicity, or stress, which has consistently been shown to condition Spanish /s/ lenition. The two conditions for this factor are *tonic* and *atonic*, with segments occurring in the former context more perceptually-salient,[11] more contextually meaningful,[12] and therefore more resistant to lenition than those occurring in the latter. I propose that rates of lenition will differ by tonicity most for PRESERVE(strident, voice), PRESERVE(coronal), and PRESERVE(segmental) as these three constraints represent the most dramatic changes to input specifications (i.e., [s] versus not [s], place of articulation versus no place of articulation, and segment retention versus deletion, respectively). I also propose that tonicity will be more important in stronger prosodic positions where cues are already more salient.

This section has presented Faithfulness constraints that are formulated according to feature(s), prosodic positions, and tonicity; a complete list of constraints examined in this dissertation and their descriptions can be found in Appendix A. While the phonetically-based

---

[11] Stressed syllables tend to be longer and louder than their counterparts; these features are a fundamental phonetic dimension of prosody.

[12] Syllable stress is contrastive in Spanish (e.g., *caso* [ˈka.so] 'I marry' versus *casó* [kaˈso] 'he/she married').

Markedness and Faithfulness constraints presented in this section aim to capture the linguistic motives for Salvadoran /s/ weakening, we must remember that this phenomenon is highly variable and therefore not predictable via strict constraint ranking in OT. For example, even if the learner knows that placing an [s] between two non-high vowels in word-final position comes at a high effort cost but a low perceptual cost, that sequence likely won't be avoided categorically, and LAZY(vcl_strid_fric, $V_{[-high]}$__$V_{[-high]}$) will not dominate the relevant lenition-blocking constraint in every tableau. As a potential solution to accounting for a grammar with free variation, the following section considers maximum entropy, a learning algorithm that uses real data to model variation in OT by assigning weights to constraints in lieu of strict ranking and assigns probabilities to each possible outcome.

3.7 A maximum entropy model of phonology

Faced with the expectation that linguistic theories should explain how a learner acquires language-specific knowledge, OT has faced one particular challenge: how to account for data with free variation. Early algorithms for learning constraint rankings reflected this issue; consistently unable to converge when fed noisy training data, they were restricted to learning a single, strict ranking. As I have alluded to in previous sections, this is problematic for phenomena such as Spanish /s/ weakening in which one input has many possible outputs that are more or less probable according the context in which that input occurs. For example, my Salvadoran data requires a model that knows that the input *casa* /kasa/ ('house') has (at least) six outputs, [ka.sa], [ka.za], [ka.$s^θ$a], [ka.ha], [ka.ɦa], [ka.a], all of which are possible but only some of which are probable given a specified set of conditions. The model should also be able to extend its knowledge about

/kasa/ to similar, novel forms that it has never seen, mimicking the process of a learner exposed to language data containing free variation.

Various attempts have been made to build a model and corresponding algorithm that can do just this: learn from real, potentially noisy data in an effort to represent both free variation and categorical distinctions and generalize to examples not in the training data. Two models have been particularly successful in these respects: Boersma's (1997) Stochastic OT model and associated Gradual Learning Algorithm (henceforth, GLA) and the Maximum Entropy model (henceforth, maxent), proposed by Goldwater and Johnson (2003). Maximum entropy is a general class of statistical models that are motivated by information theory and therefore strive to remain non-committal (i.e., high entropy) while staying faithful to the training data. Jäger and Rosenbach (2006, p. 944) define 'entropy' as a quantification of "the bias of a probability distribution" where "the least biased distribution has the highest entropy, and vice versa. If we have partial knowledge about a stochastic process and we have to estimate the underlying probability distribution, the best guess to choose among all distributions that are compatible with our knowledge is one with the highest entropy."

In their advocacy for the use of maxent over GLA, Goldwater and Johnson argue that the former is both formally and practically superior to the latter. With respect to its formal advantages, maxent is based in the principles of information theory; as a result, it is well-motivated mathematically and its properties have been well-established. Maxent also presents practical advantages in that the only parameters we must set are μ and σ^2, which specify the preferred value of each constraint and how closely the model should fit the training data to that preferred value, respectively; the effects of these parameters on the accuracy of the model are clear and undisputed. The GLA, on the other hand, requires that at least two parameters be set (the ratio of the plasticity value to the number of training examples and the evaluation noise), yet there is no clear method for choosing these parameters or knowing how sensitive the model is to these choices. Furthermore, maxent, unlike the GLA, is guaranteed to converge.[13]

Additionally, while GLA and maxent have both been lauded for their ability to learn from noisy input to build grammars that account for free variation, Stochastic OT is unable to account for the effects of cumulative constraint violations, i.e., more violations of a low-ranked constraint may outweigh fewer violations of a highly-ranked constraint, or two low-ranked constraints gang up to jointly beat a highly-ranked constraint. Indeed, in a paper comparing the accuracy of predictions made by Stochastic OT and maxent for real speech data, Jäger and Rosenbach (2006, p. 937) show that the maxent model's prediction of cumulativity, which they argue is crucial for modeling probabilistic variation in real language use, renders it the "clearly superior model, both with respect to the accuracy of its predictions and to its learnability properties." Because there is clear evidence for widespread cumulativity in the Salvadoran /s/ weakening data, as we will see in Chapter 6, a maxent approach was deemed optimal for this dissertation's analysis.

For readers unfamiliar with maxent grammars, this section has summarized the key features of the model and its utility in producing stochastic grammars. For explicit formulas, see Goldwater and Johnson (2003, pp. 112–113), Jäger and Rosenbach (2006, p. 944), and Chapter 5 of this dissertation, in which I walk the reader through maxent probability calculations using sample training data. The next and final section of this chapter addresses yet another challenge to modeling Salvadoran /s/ lenition in OT: how to account for the effects of socially-conditioned variation in the grammar. While no perfect solution has been proposed, the following section explores possible

[13] An iterative algorithm is said to converge when the changes in estimates from iteration to iteration get progressively closer to a specific value.

strategies based on accounts of stochastic grammars that incorporate non-grammatical factors into their models.

3.8 Modeling variation codetermined by social factors

While many constraint-based models for variable phenomena have been developed in recent years, the focus has been on modeling pure grammatical variation. Constraint-based grammars—developed using a variety of models including GLM, maxent, Stochastic OT, and Noisy Harmonic Grammar—typically pay little attention to non-grammatical factors, despite general agreement that these factors have an important effect on patterns of phonological variation. Indeed, in a discussion of this gap in the research, Coetzee (2016, p. 213) argues that, "Any formal model that purports to be a realistic model of the cognitive system underlying phonological variation therefore has to allow for both grammatical and non-grammatical factors to codetermine the patterns of variation."

How, then, do we build constraint-based grammars that are able to account for the interaction between grammatical and non-grammatical factors? Surely, we can learn something from variable rule analysis, first proposed by Labov (1969) to allow for the interaction between linguistic and non-linguistic factors such as communicative context and demographic characteristics of the speaker. In such analyses, traditionally implemented via logistic regression, the grammar is not privileged; that is, there is no mathematical distinction between how linguistic and non-linguistic factors affect the grammar.

Seemingly contradicting this method is the general finding by the varationist tradition that non-grammatical and grammatical factors are actually independent from one another. To exemplify this separation I turn to Labov's (1966) classic study of the social stratification of

English in New York. While Labov finds that the rate at which a speaker produces [r] in pre-consonantal and word-final positions is highly dependent on his or her social status, [r] is still significantly more likely to be produced word-finally than pre-consonantally for all speakers, regardless of social class. In this sense, both linguistic and extralinguistic factors contribute to the grammar but play different roles: grammatical factors specify possible variants and establish their frequency structure, while extralinguistic factors impose changes on this frequency structure.

To reflect this reality, I opt to build a grammar-dominant model in which linguistic and non-linguistic factors both contribute but are implemented differently. Following Coetzee (2016), grammatical factors are implemented in the form of constraints (as discussed in previous sections of this chapter) while non-grammatical factors have the unique ability to scale Faithfulness constraint weights up or down.

In his account of nasal place assimilation in English, Coetzee develops a computational model of the listener based on two perception experiments via Noisy Harmonic Grammar (Legendre et al., 1990), which models phonological variation similarly to maxent. While maxent develops a grammar by turning each candidate's harmony directly into a probability of being chosen, Noisy HG adds random noise, drawn from a normal distribution centered on zero, to each constraint's weight for every utterance. The resulting grammar on each trial chooses the winner by summing weighted violations rather than by strict ranking. Specifically, Coetzee uses a version of Noisy HG, developed by Coetzee and Kawahara (2013), that permits the weights of constraints to be scaled up or down based on non-grammatical variables. In this model, the grammar determines the space of possible variation, and non-grammatical factors determine the frequency with which specific forms are observed.

Coetzee first develops a basic grammar according to the nasal place assimilation data from his two experiments. In this model, the algorithm uses Coetzee's data to assign weights to pertinent constraints and then uses these weights to make predictions about future patterns of variation. Once this model has been created, Coetzee begins incorporating the contributions of two non-grammatical factors, inhibition and speech rate, into the model. Coetzee finds that, as a listener's inhibition increases, being unfaithful to the input is less likely than predicted by the grammar model. Consequently, Faithfulness constraints should have more weight. Contrastingly, as the speaker's speech rate increases, being unfaithful to the input is more likely than predicted by the model, and Faithfulness constraints should be given less weight.

Coetzee uses an iterative best-fit procedure to calculate the amount by which Faithfulness constraints should be scaled up or down for inhibition and speech rate: first, constraint weights were scaled up or down in increments of 0.05 for each level of the independent variable (i.e. for each speech rate). Then, output patterns predicted by each scaled version of the grammar were established using the "To output distribution…" function in *Praat*. Because output patterns varied slightly each time this function was performed due to noise in the data and the grammar, the procedure was repeated ten times and average patters were calculated over these ten iterations. The output patters predicted by this average were then compared to patterns in the observed data and the mean square error (MSE) of assimilation rates between the two models (predicted and observed) were calculated. These steps were repeated for each scaling rate; the scaling rate that produced the smallest MSE (i.e., was the best fit of the observed data) was taken as the 'scaling rate' for that given non-grammatical variable and was applied to Faithfulness constraints accordingly.

By adjusting the weights of Faithfulness constraints using these scaling rates, Coetzee creates a model in which a hearer's "Faithfulness knob" is turned up or down depending on non-grammatical factors. While grammatical factors are used to establish a baseline model for variation, non-grammatical factors work to adjust this model so that it is a better fit for real data. Along these lines, I develop a model in which all Salvadorans have the same baseline grammar that serves as a good representation of the entire speech community. However, by turning "Faithfulness knobs" up and down depending on social variables, the frequency structure of possible output forms adjusts to reflect the effect of these social variables in the real data, mimicking the experience of the learner. Section 5.6 describes this procedure, which is inspired by but not identical to Coetzee's, in more detail.

3.9 Conclusion

This chapter has explored phonetic motivations for /s/ weakening and has presented a framework for accounting for these motivations within Optimality Theory. Borrowing from Kirchner's effort-based account of consonant lenition, I have proposed a phonetically-based approach to modeling Salvadoran /s/ lenition that relies on a maximum entropy model of phonology as well as variationist theories that provide important clues for incorporating social variation into a constraint-based grammar.

The following chapter explores a specific phonetic realization of /s/ in Salvadoran Spanish in great detail in order to argue that this allophone occurs as the result of gradient lenition and should therefore be considered within the purview of this dissertation. As this proposal is new, to my knowledge, and is likely to be controversial in the field, I devote an entire chapter to exploring the acoustic and phonological nature of this sound, which has never been studied empirically in El Salvador to my knowledge.

Chapter 4. Salvadoran [s^{θ}]: Evidence for /s/ lenition

4.1 Introduction

A voiceless fricative that is perceptually akin—but not articulatorily identical—to interdental /θ/ has been attested in El Salvador, as well as in neighboring Honduras and Nicaragua (Canfield, 1981; Hualde, 2005; Lipski, 1994). While this sound has received little attention in Central America, it has been studied at length in Andalusia, where it is referred to as *ceceo* and is variously transcribed (Azcúnaga López, 2010; Hualde, 2005; Penny, 2000). Despite the body of research on Andalusian *ceceo,* the articulatory and acoustic nature of this sound, which I refer to as [s^{θ}] throughout this dissertation, has remained somewhat elusive. Penny (2000, p. 118) describes it as a dental fricative with "fronting of the tongue body so that the sound acquires some of the acoustic qualities of interdental /θ/"; Navarro Tomás (1933) argues that what differentiates it from [s] is that the tongue lays flat and unrounded. In a description that embodies both of these ideas, Canfield (1981, p. 54) describes this sound as "approach[ing] [θ], but lacking the tenseness…The impression is that of a lisp, and the tongue is flat and ungrooved."

While traditional dialectal maps suggest that regions of rural Andalusia at one point employed this allophone exclusively for orthographic 'c', 'z', and 's', this may be a linguistic idealization (Regan, forthcoming). Instead, there appears to be significant intraspeaker variation, with many *ceceante* speakers alternating between [s] and [s^{θ}] with no clear pattern (Navarro Tomás et al., 1933). Crucially, however, Andalusian *ceceo* is socially indexed, and is "considered 'rural' and evaluated negatively by Andalusian city-dwellers" (Hualde, 2005, p. 154). Today, due to a number of factors including increased access to education, there is evidence that historically *ceceante* regions are adopting the prestigious standard of *distinción,* the two-phoneme system in which the apicoalveolar /s/ corresponds to orthographic 's' and the interdental /θ/ corresponds to 'c' and 'z' (Melguizo Moreno, 2007; Moya Corral & García-Wiedemann, 1995). See Section 2.2 for a brief review of this literature.

In El Salvador, like in Andalusia, [$s^θ$] "appears to be common, but subject to individual variation" (Hualde, 2005, p. 154). Lipski (1994), in a summary of the most prominent characteristics of Salvadoran Spanish, observes its increased use in rural communities. Beyond brief allusions such as these, however, little is known about the nature of Salvadoran [$s^θ$] and why this *andalucismo* is salient in some regions of Spanish America but not others. While scholars readily acknowledge the presence of this sound in El Salvador, only Azcúnaga López (2010) has studied its use empirically to my knowledge; his chapter in Quesada Pacheco's (2010) *El español hablado en América Central: Nivel fonético* summarizes the surrounding phonological environments in which [$s^θ$] occurs. According to his findings, this allophone is preferred in syllable-onset, intervocalic position, although it's not clear if differences are statistically significant. It is, in fact, difficult to discern any patterns from his data due to a key limitation: he presents the percentage of total participants that produced a sound in a given context instead of the percentage of the time an allophone was produced in a given context for all speakers. Additionally, the small sample size (13 tokens per participant) and data collection methods limit the utility of these data as compared to those presented in this chapter.

In this same publication, five of the book's seven chapters—each of which focuses on one Central American country—document some use of [$s^θ$]. In all but Belize and Panama, this allophone occurs at least occasionally and, as with Azcúnaga Lopez's chapter, the authors provide the reader with descriptive statistics indicating the raw number and percentage of participants that produced a given allophone in a given phonological environment.[14] While this linguistic atlas, which provides an account of phonological variation in some of the field's most understudied dialects, is informative, the authors' goal is not to posit any linguistic or sociolinguistic motive for the occurrence of this sound in these dialects. As is so often the case, the authors treat this phenomenon as simply an Andalusian artifact that persists in some populations but not others.

Contrastingly, and with a focus is on Latin America more generally, Moreno Fernández (2005) endeavors a comprehensive review of work on what he calls American [θ][15] in which he explores the genesis of this phenomenon via observed patterns of variation. In this paper, Moreno Fernández divides American dialects in which this sound has been observed into two groups: those that use it systematically and those that do not. The former, which he describes as employing "false residual distinction," includes Argentina, the Canary Islands,[16] Colombia, Peru, and Paraguay. In these territories, studies such as those on Buenos Aires (Alvar, 1995), Tenerife (Alvar, 1959), and Cuzco (Caravedo, 1992) reveal that so-called [θ] is used either predominately or exclusively in contexts that correspond to phonemic /θ/ in *distincción* dialects, with rates of usage increasing in rural dialects. On the other hand, in the latter group—which includes El Salvador, Honduras, Nicaragua, and Venezuela—Moreno Fernández describes a situation in which this sound occurs sporadically within a system of *seseo*. Without citing particular studies, he adds that this type of American *ceceo* (i.e., use of so-called [θ] in free variation within an otherwise-*seseante* system)

---

[14] According to the book's introduction, all authors used identical methods for data collection and analysis. Contributing authors include Cardona Ramírez (2010a, 2010b) on Panama and Belize, Utgård (2010) on Guatemala, Hernández Torres (2010) on Honduras, Rosalis Solís (2010) on Nicaragua, and Quesada Pacheco & Vargas Vargas (2010) on Costa Rica.

[15] Moreno Fernández uses '[θ]' to refer to /θ/-like sounds generally, encompassing the numerous variations in its articulation, as described above.

[16] While the Canary Islands are not typically considered part of Spanish America and are, to this day, a Spanish territory, they were colonized concurrently with and therefore bear linguistic similarities to modern-day Latin America.

typically appears in the speech of older, rural speakers of little prestige. He attributes this to the prestige of the *seseo* norm, which dominates the media and formal academic instruction. At the end of the paper, whose stated goal is to review synchronic studies of American [θ] in an effort to better understand diachronic change, Moreno Fernández concludes that the expansion of *seseo* at the expense of *distincción* takes place in three stages, beginning with /s/-/θ/ opposition and ending in full *seseo*. It is in the intermediate stage, he asserts, that so-called [θ] becomes more sporadically-utilized in some dialects while remaining restricted to very specific lexical forms in others. Crucially, this interpretation assumes that all uses of so-called [θ] in Spanish America—regardless of the sociolinguistic distribution in which they occur—originate from the peninsular interdental phoneme. In this chapter, I challenge this assumption.

While debuccalization/deletion and Latin American *ceceo* have traditionally been treated as unrelated phenomena affecting /s/, I propose in this chapter that [s^{θ}], [h], and [Ø] are the result of varying degrees of reduction of the articulatory gestures needed to produce a sibilant, with [s^{θ}] indicating undershoot of the oral gesture, [h] indicating complete loss of the oral gesture with retention of the glottal gesture, and [Ø] indicating a loss of both oral and glottal gestures. In this sense, [s], [s^{θ}], [h] and [Ø] can be thought of as existing on an articulatory continuum in which [s] represents one extreme (strongest, most constricted) and [Ø] represents the other (weakest, least constricted).

Laying the foundation for this proposal are gradient accounts of /s/ weakening, which have examined lenition via continuous acoustic measures and are discussed thoroughly in Section 2.4. While these authors do not map instrumental measures such as duration and COG onto traditional segmental categories, doing so could lend support to the claim that these allophones are situated on an acoustic continuum, provided that values for duration and COG decrease progressively from

[s] to [s^θ] to [h] and are significantly different for each allophone. Crucially, because changes in articulatory space (which is necessarily continuous) are reflected in continuous acoustic output, we can estimate changes in the articulatory properties of a sound by measuring changes in its acoustic output. Therefore, if [s], [s^θ] and [h] exist on acoustic continua for duration and COG, one can reasonably claim that this ordering is mirrored articulatorily.

In this chapter, I will argue that [s^θ] occurs as the result of gestural undershoot of /s/ and should therefore be considered within the purview of sibilant weakening. I first motivate this claim within an effort-based approach to consonant lenition, proposing that [s^θ] represents an intermediate realization between full sibilance and debuccalization, therefore serving to reduce effort cost while preserving important perceptual cues such as [+coronal]. I then explore this hypothesis empirically via a series of experiments intended to shed light on the phonological, phonetic, and sociolinguistic characteristics of this sound in El Salvador. After discussing experimental results and their implications for the treatment of [s^θ] as a product of gradient /s/ weakening, I compare this sound acoustically to the interdental phoneme /θ/ to show that, despite their perceptual similarity and frequent conflation in the literature, [s^θ] and /θ/ are extremely disparate sounds.

4.2 Conceiving of [s^θ] as an approximant

It is not a great leap to consider [s^θ] as a product of /s/ weakening given the articulatory descriptions of [s^θ] as flat, ungrooved, and lacking tension. Such descriptions imply that the principal difference between [s^θ] and [s] is a lesser degree of precision and tenseness in the former, an idea that fits into Canfield's (1981, p. 5) narrative about *analuicismos* in the Americas:

> An examination of the *andalucismos* of a phonological character that are now part of American Spanish reveal that they are aspects of a general trend in articulation

> that one might all *lenguaplana*, a lessening of muscular tension in grooving for sibilants and in tongue-raising for alveolar and palatal articulations, all of which may correspond to an attitude of less effort.

Indeed, [s] requires a great deal of effort to produce, as discussed at length in the previous chapter. A partial failure to fully carry out the oral mechanisms required to produce an [s], then, should result in a similar but weaker sound that suffers from gestural undershoot, which Chitoran et al. (2015, p. 2) describe as occurring when "the target is not fully achieved, but rather approximated. This means that an active articulator (e.g., the tongue, the lips) begins its movement toward the target location (e.g., the teeth or the palate) but does not achieve the degree of constriction specified for it." An effort-based approach to consonant lenition provides a useful framework for contextualizing this hypothesis and subsequently testing it with acoustic data.

Within Kirchner's (2004) framework (see Section 3.5), sibilants have a high effort cost due to the tight constriction and particular precision required to produce them. Consequently, in dialects of Spanish in which speaking 'economically' is prioritized, speakers may mitigate the demands of the high-cost sibilant via debuccalization (eliminating the segment's place of articulation and therefore reducing effort cost) and/or deletion (eliminating the segment entirely). However, Kirchner points out that lenition comes at the cost of perceptual distinctions. This may, in part, explain why debuccalization and deletion are most prevalent in coda position, which is prosodically weak and less perceptually salient. A solution for dialects that seek to reduce effort cost—yet strive to maintain perceptual distinctions in strong prosodic positions—might be a lenited segment that retains some amount of oral constriction but is less effortful than a canonical [s]. I posit that the use of [s^{θ}] in Salvadoran Spanish is just this.

The following section describes how empirical data will be used to show that Salvadoran [s^{θ}] exists on acoustic and articulatory continua between [s] and [h], lending support for the proposal that it serves to satisfy competing needs: speakers are able to minimize effort cost while preserving important perceptual cues.

4.3 Methodology

In this section, I motivate and carry out a series of experiments to support my claim that Salvadoran [s$^{\theta}$] occurs as the result of /s/ lenition and is therefore within the purview of this dissertation. Given my hypotheses regarding the phonetic origins and articulatory nature of this sound, I propose that [s$^{\theta}$] has unique utility within Kirchner's effort-based framework, as it achieves a lower effort cost than [s] by undershooting the oral gesture yet retains more perceptual salience than debuccalized [h]. Table 7, below, summarizes the hypotheses I have set forth thus far by comparing [s], [s$^{\theta}$], and [h] along three dimensions: status of the oral gesture, relative effort cost, and perceptual salience.

Table 7

Proposed continuum, [s] to [h]

Dimension	[s]	[s$^{\theta}$]	[h]
Oral gesture	Fully realized	Undershot	Eliminated
Relative effort Cost	High	Medium	Low
Perceptual salience	High	Medium	Low

The following section briefly summarizes how the data used in the experiments presented in this chapter were collected and analyzed. For a complete discussion of methods for data collection and acoustic analysis, see Chapter 5.

4.3.1 Data collection and analysis

Salvadoran participants for experiments presented in this chapter are the same 72 speakers whose speech data is used to develop the maxent model in Chapter 6 of this dissertation. Chapter 5 provides a detailed account of participant recruitment and criteria, recording equipment and procedure, segmentation, and acoustic analysis that will not be repeated here. Chapter 5 also provides spectrographic comparisons of the sounds that are the focus of the present chapter (Salvadoran [s], [s^{θ}], and [h]), which may be of interest to the reader. A crucial difference in the procedure used for the analyses presented in this chapter and those presented in Chapter 6 is that the dataset is half as large: while this chapter takes the first 100 occurrences of phonological /s/ per participant (n = 7,200 tokens) as its object of analysis, the maxent model uses the first 200 occurrences (n = 14,400 tokens).

While a total of 7,200 tokens of phonological /s/ were coded for the purpose of the forthcoming experiments, the acoustic analyses presented in Section 4.4.1 focus exclusively on phonetic realizations categorized as [s], [s^{θ}], and [h],[17] of which there were only 5,468 tokens. For these 5,468 tokens, duration and mean COG measures were subsequently extracted using a script developed by Lars Hinrichs (Department of English, University of Texas at Austin) that measures COG at three points: one quarter, one half, and three quarters of the way through the segment. Following Erker (p.c.), I averaged these three measures to obtain a mean COG value for each token. Additionally, following File-Muriel and Earl Brown (2010), frequencies below 750 Hz were excluded from COG measurement via a Hann low pass-band filter in an effort to capture the noise component of the fricative while excluding glottal pulsing.

[17] For the purposes of this chapter, whose focus is degrees of oral constriction, tokens of [s] and [z] and tokens of [h] and [ɦ] have been collapsed.

To obtain acoustic data on phonemic /θ/ for the cross-dialectal comparison, four native speakers of Standard Peninsular Spanish—a dialect that contrasts /s/ and /θ/ phonemically—were recorded using the same recording equipment and procedure as described in Chapter 5. Participants had been living in Los Angeles for less than five years and were balanced for gender. Beginning ten minutes into each recording, the first ten tokens of intervocalic, phonemic /θ/ per speaker (n = 40) were segmented and coded in *Praat*, and duration and COG measures were extracted using the methods described above.

4.4 Results

4.4.1 Acoustic evidence for lenition

While characteristics of [s^{θ}] along the three dimensions presented in Table 7 are hypothesized, the articulatory and acoustic nature of both [s] and [h] are well-established. Therefore, if we know the effect that particular articulatory differences have on specific acoustic correlates—in this case, duration and COG—we can use measures of those acoustic correlates for [s] and [h] as anchors on a continuum and test whether measures for [s^{θ}] fall in between those anchors. If [s^{θ}] does, indeed, fall between [s] and [h] for measures of duration and COG, that ordering can then be mapped onto corresponding articulatory properties as partial evidence for gestural undershoot. With this goal in mind, the following sections present acoustic evidence that supports the continuum proposed in Table 7.

4.4.1.1 Duration

In order to ensure that the data were normally distributed, I centered and scaled duration values for each of the three allophones (n = 5,468 tokens), excluding any observations more than two standard deviations from the respective mean. In total, 182 observations were excluded: 126 tokens of [s] (3%), 14 tokens of [s^{θ}] (4%), and 42 tokens of [h] (4%), leaving 5,286 observations to be included in the duration plots and analyses below.[18]

To determine whether [s], [s^{θ}], and [h] exist on a temporal continuum, means and distributions for duration values were compared graphically in R (R Core Team, 2017) using the ggplot2 package (Wickham, 2009). All graphs presented in this chapter were created using ggplot2 unless otherwise noted. To verify that duration measures for each sound were statistically different from one another, a linear model was then constructed.

A graphical comparison of duration measures for [s], [s^{θ}], and [h] reveals that the three sounds do, indeed, exist on a continuum, with tokens of [s] generally the longest, followed by [s^{θ}] and then by [h]. Figure 14 shows this distribution.

[18] It should be noted that the purpose of scaling was solely to identify outliers; duration values that appear in plots and statistical analyses are based on raw measures in milliseconds.

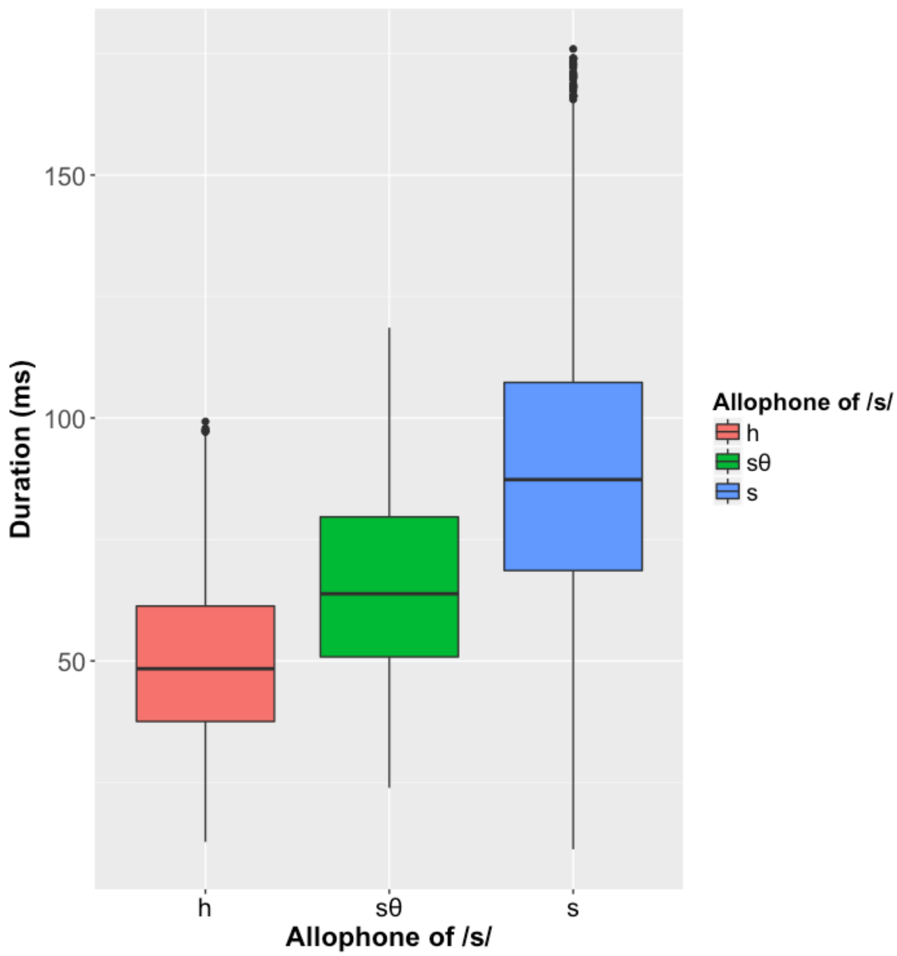

Figure 14. Duration of /s/ (in ms) by allophone, from [h] (shortest) to [s] (longest).

If [$s^{θ}$] is produced with some amount of gestural undershoot as compared to [s], as hypothesized, it should require less time to produce than the gesturally-complex strident requiring sustained antagonistic force, but more time to produce than [h], for which there is no oral constriction. This ordering is reflected in Figure 14. While [s] and [h] show no overlapping duration values between the first and third quartiles of the data, measurements for [$s^{θ}$] appear to fall right in between the two, with measures between the first quartile and median overlapping substantially with measures for [h] and those between the median and third quartile overlapping with measures for [s].

Indeed, the linear regression model revealed that duration measures for all three sounds are statistically different from one another. Table 8 summarizes these results.[19]

Table 8

Linear regression for effect of allophone on segment duration in milliseconds

Allophone	β	Standard Error	*t* value	*p* value
Intercept ([s^{θ}])	66.55	1.65	40.35	<0.001***
[s]	21.87	1.71	-12.79	<0.001***
[h]	-15.99	1.88	-8.49	<0.001***

These results show that the allophones [s], [s^{θ}], and [h] each occupy their own unique space on a duration continuum. When mapped onto articulation, this suggests that [s] requires significantly more time on average to produce than [s^{θ}], which in turn requires significantly more time to produce than [h]. These findings represent initial evidence for the claim that [s^{θ}] occurs as the result of gestural undershoot of /s/ and can therefore be situated within the phenomenon of gradient consonant lenition. The following section presents a second piece of evidence in support of this claim.

4.4.1.2 Spectral COG

An identical procedure was used to examine spectral COG. As with duration, COG measures outside two standard deviations for each allophone were excluded. In total, 279 observations were excluded: 213 tokens of [s] (5%), 16 tokens of [s^{θ}] (5%), and 50 tokens of [h]

[19] To address concerns regarding the effects of coarticulation and onset/coda asymmetries on duration values, identical analyses were run on a subset of the data in which all tokens occurred in syllable-onset, intervocalic position (n = 2,418). As expected, durations were slightly longer for all allophones, with means for tokens of [s] increasing from 88.42 to 93.49 ms, those for [s^{θ}] increasing from 66.55 to 74.88 ms, and those for [h] increasing from 50.56 to 56.71 ms. Differences by allophone were still statistically significant (p<0.001 for all comparisons).

(5%), leaving 5,190 observations to be included in the COG plots and analyses below. To determine whether [s], [s^{θ}], and [h] exist on a continuum with respect to COG, means and distributions for COG values were compared graphically. Just as with duration, I verified that COG measures for each allophone were statistically different from one another by constructing a linear model in R.

A graphic comparison of COG measures for [s], [s^{θ}], and [h] reveal that the three sounds also exist on an acoustic continuum for this measure, with tokens of [s] generally the highest, followed by [s^{θ}] and then [h]. Figure 15 shows this distribution.

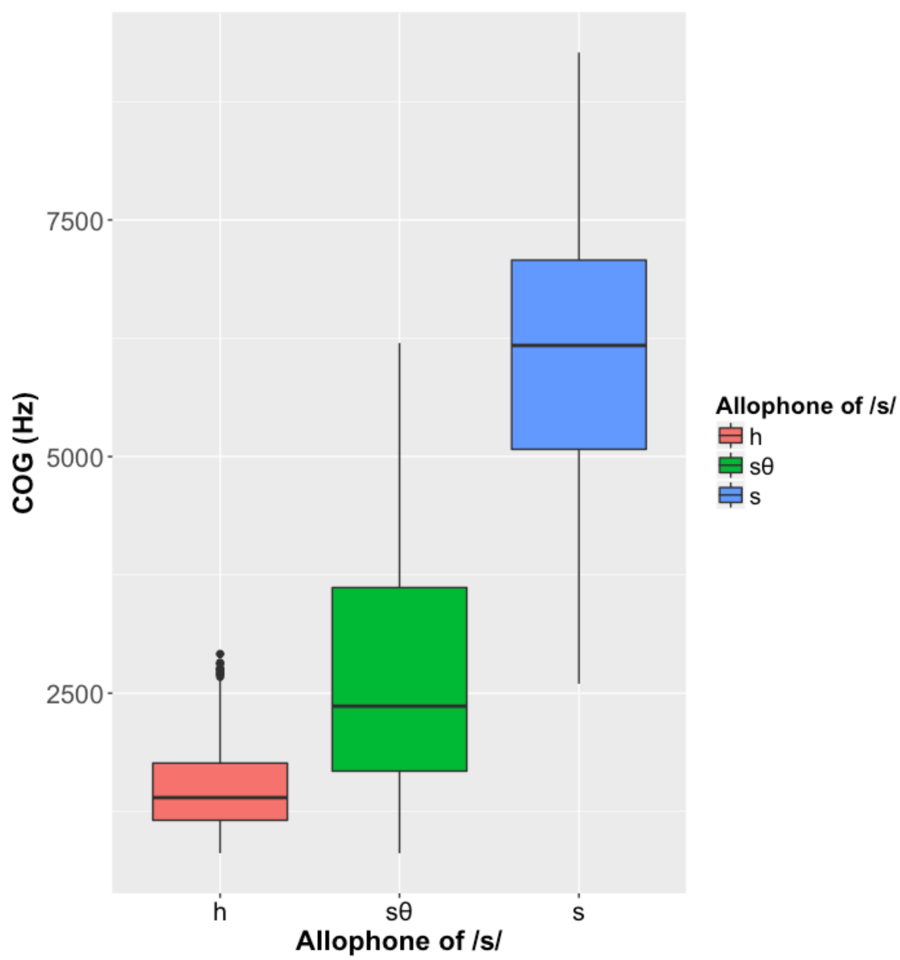

Figure 15. COG of /s/ by allophone, from [h] (lowest) to [s] (highest).

Here we see that COG measures, dependent on both frequency and amplitude, are notably lower for [s^θ] than for [s]. This supports the hypothesis that [s^θ] is produced with some amount of gestural undershoot: when constriction at the place of articulation is relaxed, airflow turbulence—which generates noise at high frequencies—decreases. Using COG as a tool to approximate the vocal tract's configuration when producing each of these sounds, we see that [s^θ] occupies the space between tightly-constricted [s] and debuccalized [h].

Unsurprisingly, given the differences observed in Figure 15, the linear regression model revealed that COG measures for all three sounds are statistically different from one another.

Table 9

Linear regression for effect of allophone on spectral COG in Hertz

Allophone	β	Standard Error	*t* value	*p* value
Intercept ([s^θ])	2774.39	76.37	36.33	<0.001***
[s]	3262.35	79.21	41.18	<0.001***
[h]	-1267.03	87.27	-14.52	<0.001***

Similar to the conclusions drawn from the duration comparison, these results show that each of the three allophones occupies different acoustic space. However, the differences are even starker in terms of COG, a measure for which there is virtually no overlap between the first and third quartiles for any pair of allophones. The observation that these three sounds are better defined by differences in COG than by differences in duration can be quantified by comparing the R-squared values for the two models: while the duration model only accounts for 22% of variation in the data (r^2=0.22), the COG model accounts for 66% (r^2=0.66), revealing that the data are three-times better fitted to the regression line in the latter. Additionally, both Figure 15 and the results of the logistic regression test reveal that in terms of COG—a good proxy for the degree of constriction in the oral cavity—[$s^θ$] is more similar to [h] than to [s]. This confirms the hypothesis that [$s^θ$] does require significantly less constriction than [s] and, in turn, has a significantly lower effort cost within Kirchner's framework.

4.4.1.3 Summary: Acoustic evidence

In this section, I have operationalized /s/ weakening as both the shortening of duration and the lowering of COG. Figure 16 shows the relationship between these two continuous variables in the data as well as the transcribed allophone assigned to each data point.

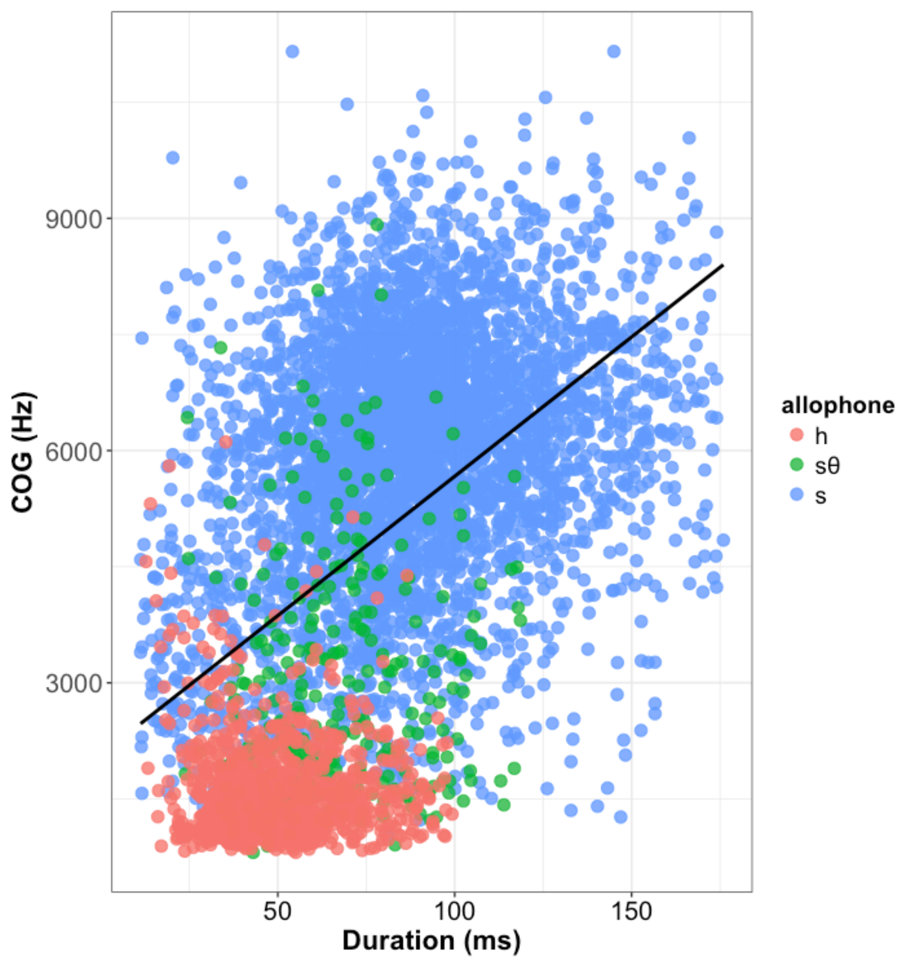

Figure 16. Relationship between duration (ms) and COG (Hz) by transcribed allophone.

Here, we see that the correlation between COG and duration is positive and moderately strong (r=.50). Moreover, we see that tokens of [h] cluster in the bottom left of the graph, tokens of [s] cluster in the top center and right, and that tokens of [s^{θ}] are scattered in between these clusters. This figure summarizes the findings of this section, in which I have shown that [s^{θ}] exists on an acoustic (and therefore articulatory) continuum between [s] and [h], lending support to the proposal that it occurs as the result of gestural undershoot, which serves to ease articulatory effort cost while preserving important perceptual distinctions.

Furthermore, this section has argued that the distinct groupings of duration and COG values for these three allophones (which is somewhat obscured in Figure 16) suggest that these sounds occupy their own acoustic spaces; while there is some overlap in these spaces, they remain significantly different from one another. This distribution is best examined via a density plot, which visualizes the distribution of data over a continuous interval and shows where values are concentrated throughout that interval. Figure 17a and Figure 17b shows density plots for the measures of duration and COG in the data, respectively.

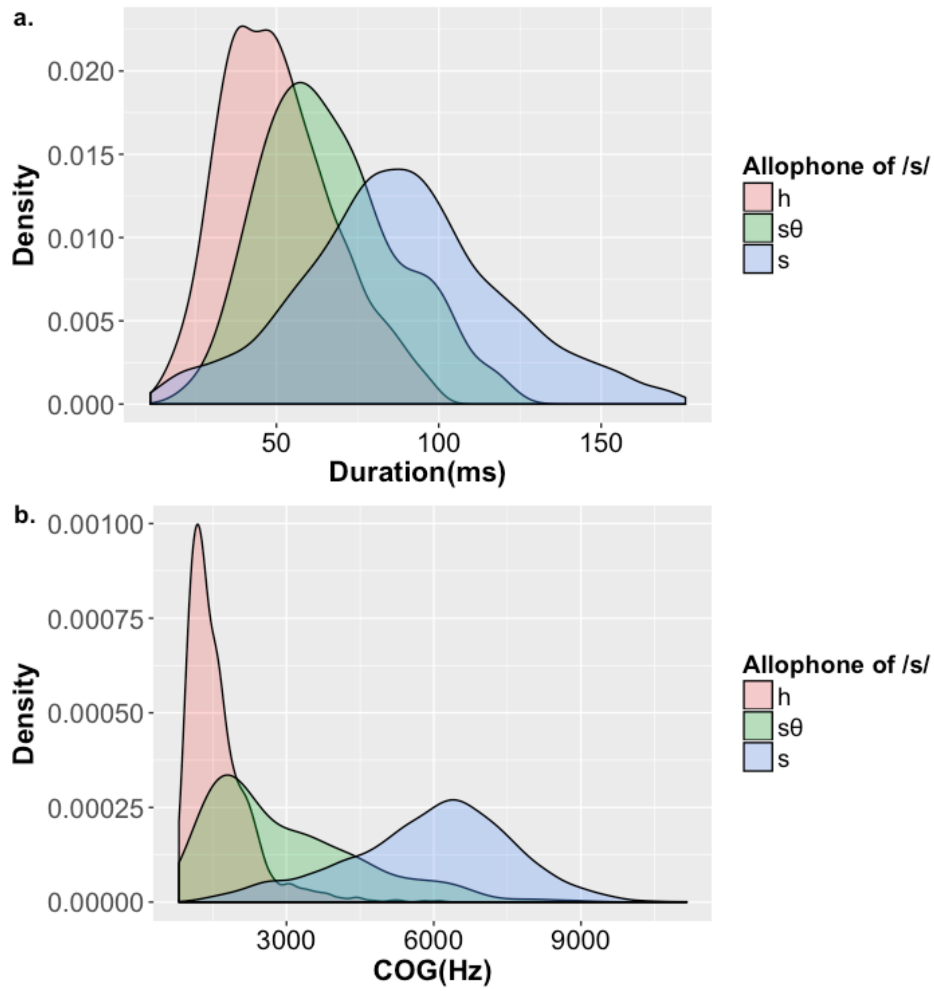

Figure 17. Density plots showing three distinct concentrations of duration (a.) and COG (b.) measures along a continuous scale.

Evident in both the scatter and density plots is the appearance that [s] and [h] are phonetic categories that a speaker can aim for, while [s^θ] occurs at their edges, likely as the result of an undershot articulatory target.

This section has provided acoustic evidence for considering [s^θ] as the result of lenition. In the following section, I identify the linguistic environments that favor use of [s^θ] and show that these patterns lend further evidence to the claim that this allophone occurs as the product of gradient /s/ weakening.

4.4.2 Evidence from linguistic environment

Within the framework I have set forth, if [s^{θ}] serves to ease articulatory effort cost while preserving perceptual distinctions, it should occur in perceptually salient positions in which [h] and [Ø] are disfavored, yet in which [s] may also be slightly disfavored due to increased effort cost as a result of the surrounding environment. Thus, in this section I examine both the prosodic positions and surrounding phonological environments in which [s^{θ}] is most often employed. While much is known about the prosodic positions and phonological environments that favor debuccalization and deletion, this section is interested in determining whether [s^{θ}] occurs in similar contexts or in complementary distribution. As described in detail in Chapter 2, the three factors examined in this analysis—*prosodic position*, *preceding segment,* and *following segment*—have consistently been identified as good predictors for patterns of /s/ weakening. The following sections briefly describe how I conceive of each variable in my analysis and how this conception is crucial to answering the question at hand.

4.4.2.1 Prosodic position

To understand the effects of prosodic position on consonant weakening, it is crucial to understand the relationship between prosody and the phonetic encoding of segments, which I describe in detail in Section 3.6.2. The basic idea is this: the phonetic realization of an individual speech segment depends, in part, on its position within the larger prosodic structure of the intonational phrase in which it occurs, with stronger segments preferring stronger prosodic positions. By taking into account the prosodic position in which a given /s/ occurs, then, we can identify that context as prosodically strong (i.e., prominent) or weak relative to other contexts.

In the present analysis, I adopt Keating's (2006, p. 171) proposal of domain-initial effects, also described in Section 3.6.2. Within this framework, we can make some predictions about the prosodic positions that favor stronger segments and should therefore favor [s] and disfavor [h] and [Ø]. If [s^{θ}] is, indeed, a solution for maintaining prosodic strength relations (stronger segments in more prominent positions) while reducing effort cost, we would expect to see this sound primarily in quasi-strong prosodic positions (that is, in word- and syllable-initial positions, but not phrase-initially). To examine these potential differences, I have subcategorized the factor *prosodic position* into four distinct levels: phrase-initial, word-initial, syllable-initial, and syllable-final (encompassing all coda /s/s, both word-medial and word-final).[20]

4.4.2.2 Preceding segment

As I discuss in Chapter 2, the segment preceding /s/ has been shown to condition /s/ weakening when the segment occurs in syllable-initial position. In their study of Chihuahuan Spanish, Esther Brown and Torres Cacoullos (2002, 2003) find that preceding non-high vowels favor weakening significantly more than their high counterparts, and that /s/ is retained most frequently after a preceding consonant. Both Esther Brown (2005) and Earl Brown and Esther Brown (2012) find identical results in New Mexico and Cali, Colombia, respectively. These findings align with Kirchner's (2004, p. 323) assertion that more open flanking segments favor consonant lenition because the main articulator must move further to reach its target when its starts from further away, therefore increasing biomechanical effort cost. That is, non-high vowels require

[20] See Section 3.6.2 for details on how positions are defined in this dissertation.

the mouth to be the most open, with the tongue furthest from the alveolar ridge; high vowels still require no oral constriction, but the jaw is close to the top of the mouth, facilitating the movement of the tongue for a following /s/; and consonants require some degree of oral constriction, making the transition from a consonant to a following /s/ the least effortful. Following these studies, I assign levels to the variable of *preceding segment* as follows, here ordered from easiest (favoring retention) to most effortful (favoring weakening): consonant > high vowel > non-high vowel. If my proposal regarding the nature and utility of [s^{θ}] is correct, this allophone should appear least frequently after consonants and most frequently after non-high vowels.

4.4.2.3 Following segment

As discussed in Chapter 2, traditional accounts of Spanish /s/ lenition have frequently identified the phonological segment following /s/ as a crucial predictor of weakening. Often, such accounts are particularly (or exclusively) interested in the nature of the consonant following /s/ (be it the manner/place of articulation, voicing, or another articulatory differentiator) because they are looking at coda /s/ weakening. However, I have already hypothesized that [s^{θ}] will seldom occur in syllable-final contexts, as this position is both prosodically weak and articulatorily challenging, making it a prime site for debuccalization and deletion. Therefore, I am not concerned with the nature of the following consonant as much as the presence of a consonantal sound in general. As with *preceding segment*, my main preoccupation is the openness of the flanking segment; thus, I have subcategorized the variable *following segment* into the same levels as above with the addition of *following pause* to account for domain-final tokens. Again, I hypothesize that [s^{θ}] will occur most readily before a non-high vowel.

4.4.2.4 Logistic model

To determine the linguistic environments that favor [s^θ], I constructed a binary logistic regression in R that identifies significant predictors of [s^θ] versus other allophones of /s/. Therefore, all 7,200 tokens of /s/ were used to build this model. A random effect of 'participant' was also included, as the 72 participants recruited for this dissertation are intended to represent the larger Salvadoran population. Table 10 summarizes the results of this model.

Table 10

Logistic regression model comparing outcome of [s^θ] to all other allophones

	β	Standard error	*p* value
Intercept	-4.95	0.45	<0.001***
Prosodic position			<0.001***
Phrase-initial	Ref	-	-
Word-initial	0.25	0.35	0.48
Syllable-initial	0.75	0.33	<0.05*
Syllable-final	-1.71	0.43	<0.001***
Preceding segment			0.20
Consonant	Ref	-	-
High vowel	0.24	0.28	0.39
Non-high vowel	0.36	0.21	0.08
Following segment			0.76
Consonant	Ref	-	-
Pause	0.19	0.45	0.67
High vowel	0.31	0.26	0.24
Non-high vowel	0.55	0.23	<0.05*

The strongest predictor of use of [s^θ] is prosodic position, with this allophone highly favored in syllable-initial position and disfavored in both phrase-initial and syllable-final positions. These findings are consistent with predictions that [s^θ] should appear most frequently in quasi-strong positions. These patterns lend further support to the argument that [s^{θ}] occurs as the result of gradient /s/ weakening and serves as an optimal realization in salient prosodic positions in which debuccalization and deletion come at a high perceptual cost. The proportional use of [s^{θ}] by prosodic position can be visualized in Figure 18.

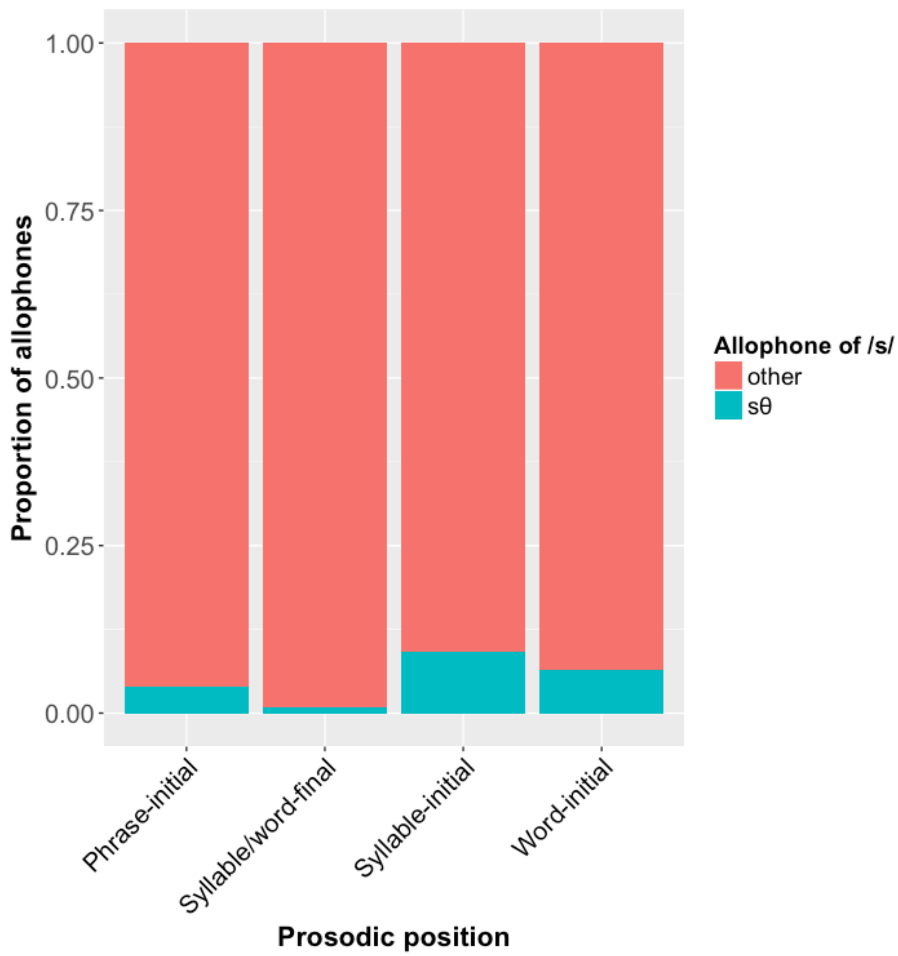

Figure 18. Proportional use of [s^{θ}] by prosodic position showing that this allophone is favored in syllable-initial (word-medial) contexts.

While the identity of the segment preceding /s/ did not reach the level of significance in this model, we see that—consistent with my hypotheses—[s^{θ}] occurs most frequently after a non-high vowel. Because this context occurs approximately seven times more than the second most-common context in the data set, this level of the variable only trends toward significance in the model (p=0.08). While this issue might be less pronounced in a larger data set, the preferred CV syllable structure of Spanish as well as the separation of vowel groups by height[21] would likely render a similar challenge regardless.

Finally, the identity of the segment following /s/ was the least significant factor in the model (p=0.76). This is unsurprising, as previous studies on Spanish /s/ weakening have shown that the preceding segment is most important in determining patterns of lenition for onset /s/ while following segment matters most with respect to coda /s/. (Because Salvadoran [s^{θ}] is primarily a syllable-onset phenomenon, we expect that the preceding segment will have a stronger effect on lenition patterns). However, while the following segment is the least important environmental factor overall, the effect of a following non-high vowel is larger than that of a preceding high vowel and reaches a level of statistical significance (p<0.05). Figure 19, which shows the proportional use of [s^{θ}] by following segment, reveals patterns similar to those observed with respect to preceding segment: the more open the flanking segment, the more likely it is that [s^{θ}] will surface.

[21] Spanish has three non-high vowels (/e, a, o/) but only two high vowels (/i, u/).

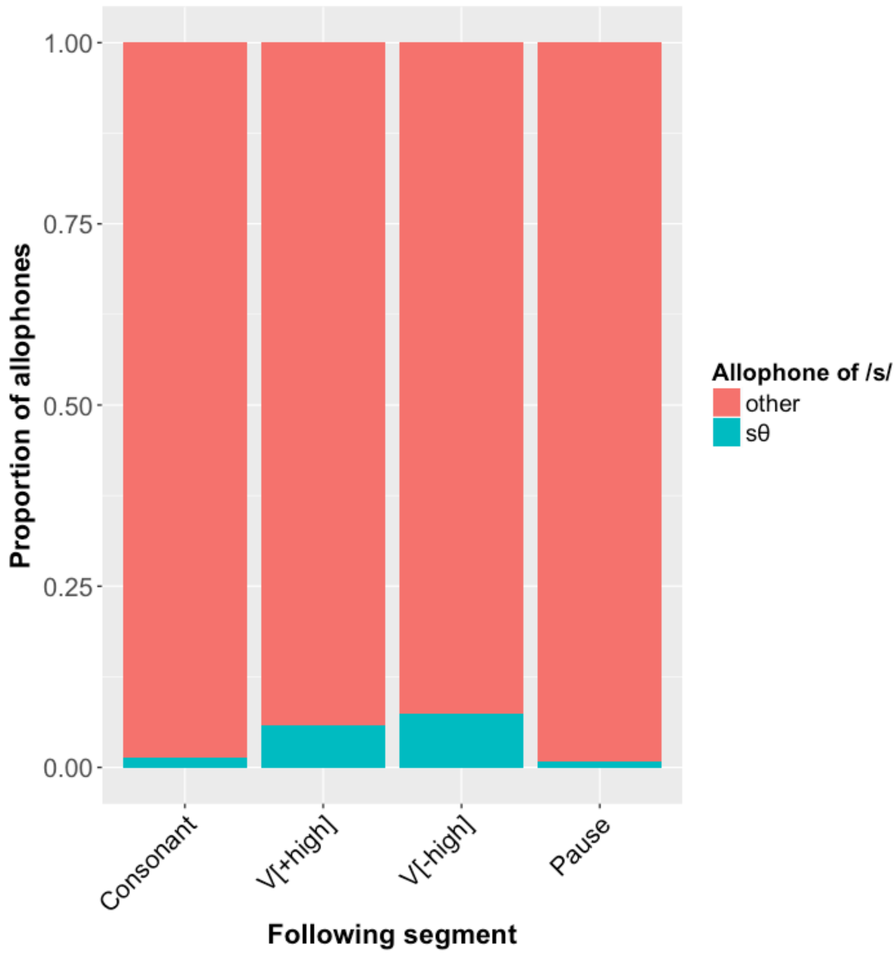

Figure 19. Proportional use of [s^{θ}] by following segment showing that this allophone is favored in contexts in which it precedes a non-high vowel.

4.4.2.5 Summary: Evidence from linguistic environment

In this section, I have provided evidence from linguistic environment to support my claim that [s^{θ}] occurs as the result of gradient /s/ weakening driven by a need to reduce articulatory effort cost while maintaining important perceptual distinctions. Specifically, I have shown that, the more open the segments flanking /s/, the more likely [s^{θ}] is to occur. Because the energy needed to move the tongue to and from the alveolar ridge increases when adjacent segments are more open, these patterns can be attributed to a desire to mitigate articulatory effort cost within Kirchner's framework. Furthermore, I have shown that [s^{θ}] occurs most frequently in prosodic positions that are perceptually salient enough to disfavor [h] and [Ø], but not so salient that they require use of [s]. Taken together, these results suggest that Salvadoran [s^{θ}] may occur as the result of competing constraints: the environment in which /s/ occurs determines the need to reduce effort cost, but only in prosodic positions in which perceptual distinctions will not be compromised.

4.4.3 Comparison to interdental /θ/

The goal of this chapter is to justify the inclusion of the [s^{θ}] allophone in my account of Salvadoran /s/ weakening. Because my assertion that this sound is the result of socially-conditioned lenition is, to my knowledge, a new proposal, I feel that an important part of my task must be to address the shortcomings of previous theories and treatments of this sound in the literature. As I briefly discuss in the introduction to this chapter, allophonic [s^{θ}] has often been conflated with phonemic /θ/; while scholars are quick to recognize that the two sounds are not articulatorily identical, papers such as that of Moreno Fernández (2005) argue that this sound exists in Latin America as a remnant of Peninsular distinction, or that it's otherwise related to the interdental phoneme. While it is impossible to rule out these theories entirely, I believe, for the reasons presented in this chapter, that a historical explanation alone is insufficient to understand the use of [s^{θ}] in El Salvador.

As dialectologists have acknowledged, the perceived fronting of the tongue body in the articulation of [s^{θ}] results in perceptual similarity with /θ/, which I believe contributes to the need to draw a connection between these two sounds. In this section, I present an acoustic comparison of Salvadoran [s^{θ}] with Peninsular /θ/ to demonstrate how disparate these sounds are with respect to production. While some regions of Latin America—particularly those that existed within the

European sphere of influence for an extended period of time—may show evidence of residual false distinction, the goal of this acoustic comparison is to show that Salvadoran [s^θ] and /θ/ have little in common beyond some perceptual similarity.

For the reader's reference, Figure 20 shows a sample spectrogram of the word *parecía* /paɾeθia/ ('it seemed'), spoken by a female participant from Madrid, which can be compared the Salvadoran token of [s^θ] presented in Section 5.3, Figure 31.

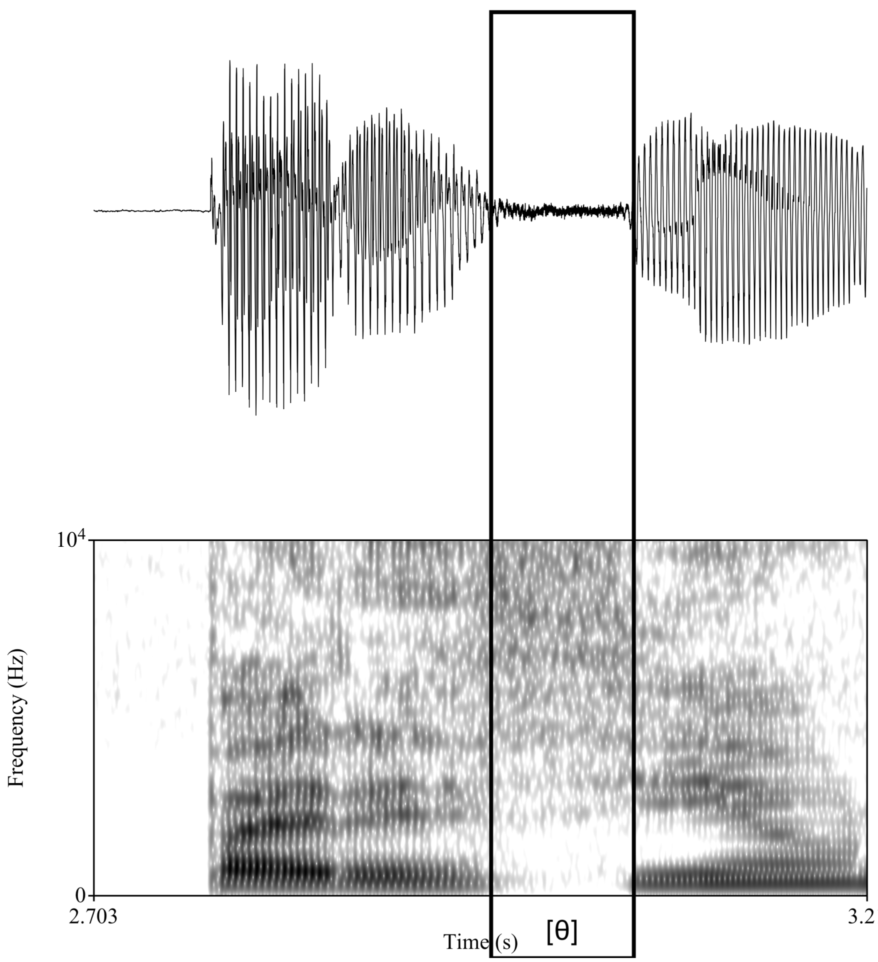

Figure 20. Waveform and spectrogram of [θ] in parecía, [pa.ɾeˈθi.a] ('it seemed'), spoken by a female participant from Madrid.

To provide a quantitative acoustic comparison of [s^{θ}] with phonemic /θ/, 40 tokens of /θ/ produced by the speakers of the Standard Peninsular dialect were compared to 40 randomly-selected tokens of intervocalic [s^{θ}] from the Salvadoran data, also balanced for speaker gender. Figure 21 shows duration measures for these two sounds.

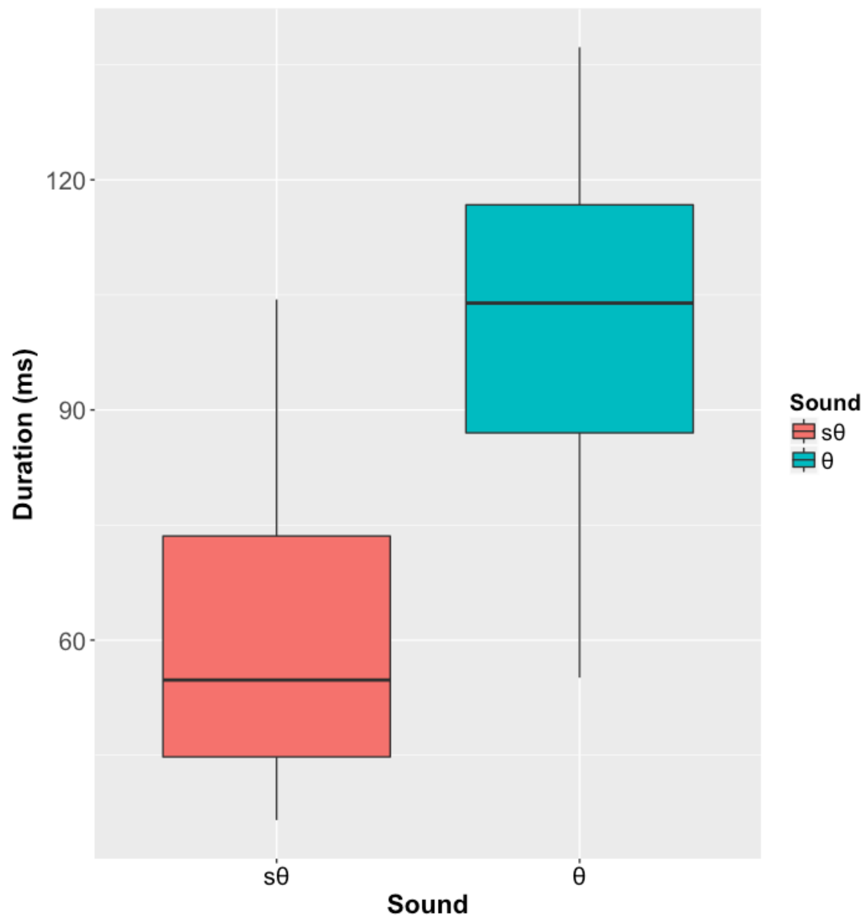

Figure 21. Comparison of duration measures for Salvadoran [s^{θ}] and Standard Peninsular /θ/.

Figure 21 shows that average durations for /θ/ and [s^{θ}] are disparate, with the former sound almost twice as long as the latter. Because all tokens occurred intervocalically, this is likely not an effect of prosodic position or surrounding phonological environment. A similar disparity was found with respect to COG measures, which are visualized in Figure 22.

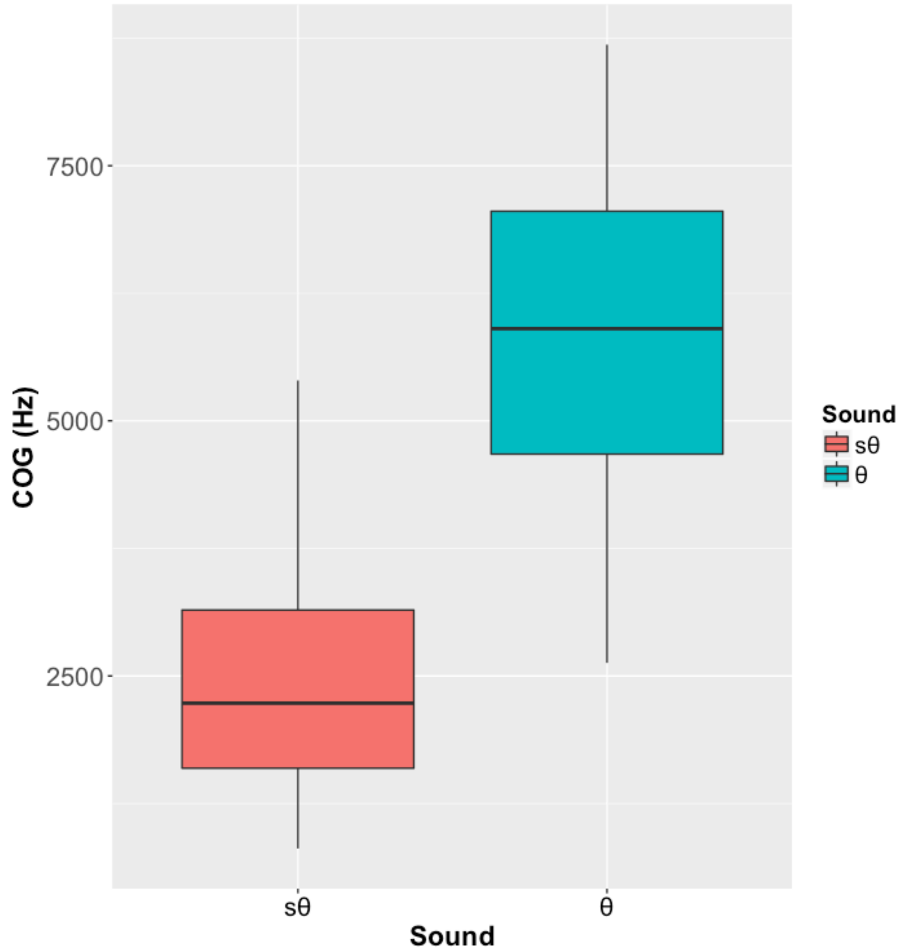

Figure 22. Comparison of mean COG measures for Salvadoran [$s^{θ}$] and Standard Peninsular /θ/.

These data reveal that, despite perceptual similarity, /θ/ and [$s^{θ}$] are extremely different in terms of production. These dissimilarities in acoustic properties reveal that [$s^{θ}$] is much weaker, both temporally and spectrally, than the interdental phoneme.

In order to further emphasize the disparate acoustic nature of these sounds, Figure 23 and Figure 24 compare measures for duration and COG across all four sounds—Peninsular /θ/ and Salvadoran [s], [$s^{θ}$], and [h]—to reveal that [$s^{θ}$] is, indeed, more similar to both Salvadoran [s] and [h] than it is to /θ/.

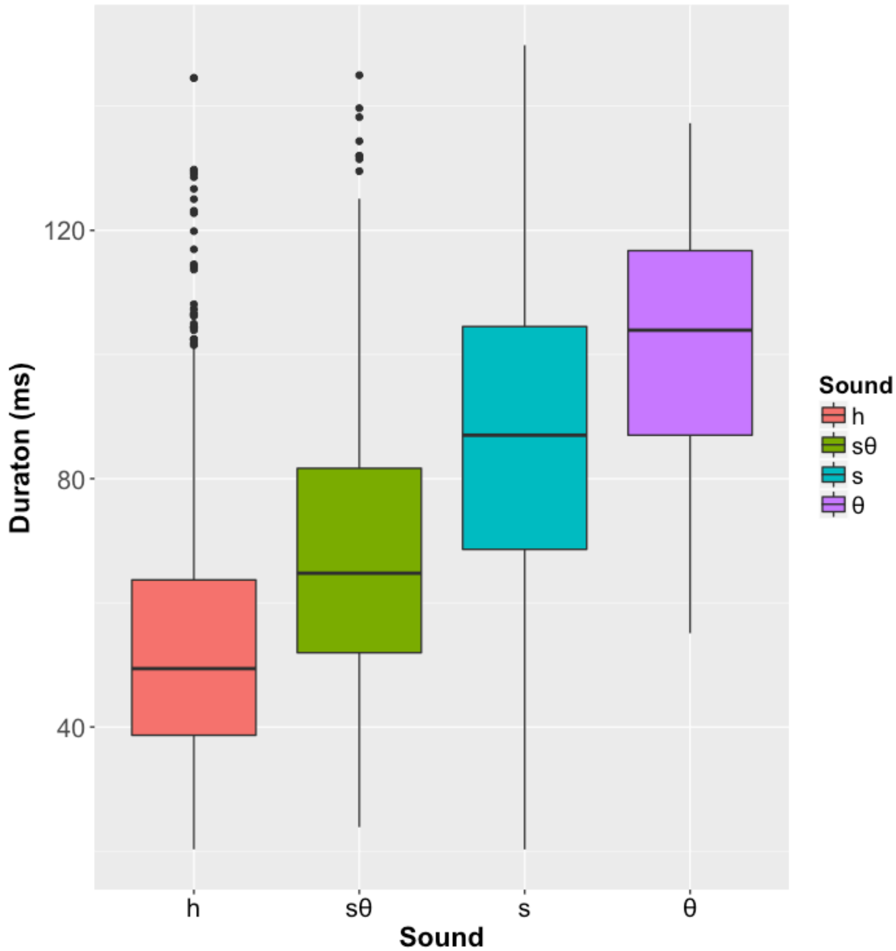

Figure 23. Comparison of duration measures for Peninsular /θ/ and Salvadoran [s], [s^{θ}], and [h].

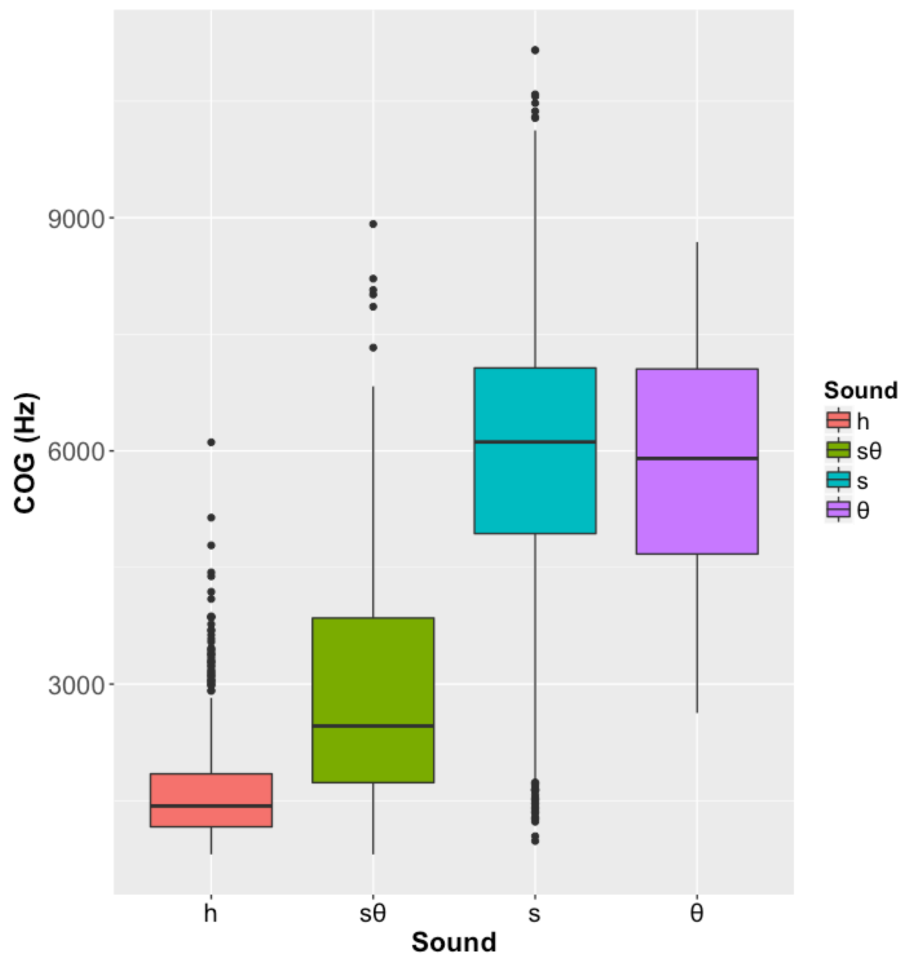

Figure 24. Comparison of COG measures for Peninsular /θ/ and Salvadoran [s], [s^{θ}], and [h].

In Figure 23 and Figure 24, we see that Salvadoran [s^{θ}] is in fact more acoustically similar to [s] than it is to /θ/. Indeed, while dialectologists have failed to go the extra step and posit a phonological connection between predorsal (Latin American) [s] and [s^{θ}], articulatory similarities have not been lost on them. Moreno Fernández (2005, p. 61; my translation) writes:

> The close relationship between the predorsal *s* and the *ceceo* phenomenon was already highlighted in 1933 by Navarro Tomás, Espinosa and Rodríguez Castellano in their famous work on the border of Andalusia (Navarro Tomás, 1975, p. 61). There they affirmed that it is significant that the Andalusian area in which predorsal *s* is used coincides with that in which *ceceo* is most common, a coincidence that makes sense when one knows the articulatory proximity of *s* and [so-called] θ in these geographical areas.

As dialectologists like Navarro Tomás point out, it's likely not coincidental that Spanish's predorsal [s] and *ceceo,* which claim the same place of origin, are articulatorily proximate. What's also not coincidental, in my view, is that Andalusian Spanish is an /s/-leniting dialect—as are all the varieties of Latin American Spanish in which some form of *ceceo* has been observed.

4.5 Conclusion

This chapter has argued that Salvadoran [s^{θ}] occurs as the result of gestural undershoot of /s/ and can therefore be considered a product of sibilant weakening within an effort-based theory of consonant lenition. In order to demonstrate exactly how each piece of evidence presented in this chapter serves to support this hypothesis, I return to Table 7, which has been replicated below for convenience.

Table 7
Proposed Continuum, [s] to [h]

Dimension	[s]	[s^{θ}]	[h]
Oral gesture	Fully realized	Undershot	Eliminated
Relative effort Cost	High	Medium	Low
Perceptual Salience	High	Medium	Low

Table 7 proposes a continuum from [s] to [h] along three dimensions: oral gesture, relative effort cost, and perceptual salience. I hypothesize that, due to the need for Salvadoran speakers to minimize effort cost while preserving important perceptual distinctions, an undershot allophone (i.e., an approximant) of [s] surfaces. While the exact nature of the oral gesture of [s^{θ}] is impossible to ascertain without electropalatographic or ultrasound technology, the experiments presented in this chapter have provided crucial acoustic and environmental information about the sounds on this continuum that work in tandem to support this hypothesis. Table 11 shows the same continuum as above but with alternate dimensions based on the experiments presented in this chapter:

Table 11

Proposed continuum, [s] to [h] with alternate dimensions

Dimension	[s]	[s$^{\theta}$]	[h]
Mean duration (ms)	92.43	69.63	53.60
Mean COG (Hz)	5932.13	2995.16	1614.46
Preferred prosodic position[22]	phrase-initial; word-initial	syllable-initial (word-medial)	syllable-final
Preferred environments[23]	C__; __$V_{[+high]}$	$V_{[-high]}$ __; __$V_{[-high]}$	$V_{[-high]}$ __; __$V_{[-high]}$

Table 11 summarizes the key findings of this chapter and demonstrates how these findings support the hypothesis that [s], [s$^{\theta}$], and [h] exist on an /s/ weakening continuum. The first two dimensions show that these sounds exist on an acoustic continuum with respect to two measures, frication duration and spectral COG, which have been established as good correlates for weakening. These measures speak to the nature of the oral gesture of [s$^{\theta}$], revealing that it requires less oral constriction and time to produce than [s], and more than [h]. From this conclusion, we can further posit that [s$^{\theta}$] represents a reduction in effort cost for the speaker as compared to [s]; if [s$^{\theta}$] is the result of gestural undershoot, as this data suggests, the biomechanical energy needed to move the tongue to the target is reduced because the target is never fully achieved. This theory is similarly supported by the *preferred environments* dimension, which shows that [s] prefers

[22] Defined as the prosodic position in which the allophone occurred the most frequently in terms of a percent.

[23] Defined as the surrounding phonological environments in which the allophone occurred the most frequently in terms of a percent.

flanking segments with more oral constriction (a preceding consonant and a following high vowel). As Kirchner (2004) asserts, rates of lenition depend in part on the openness of adjacent segments; more open flanking segments increase the effort needed to move the main articulator to its target, whereas more constricted flanking segments reduce effort cost.

With respect to perceptual salience, Table 11 shows that these three sounds also exist on a continuum determined by prosodic strength relations as defined by Keating (2006, p. 169). Confirming my original hypothesis, I find that [s] is more likely to occur the stronger the prosodic position; that is, it is most common in phrase-initial position, followed by word-initial position, followed by syllable-initial (word-medial) position, followed by syllable-final position. This reflects the phenomenon of domain-initial strengthening in which phonetic segments are stronger when they occur at the left-most edge of one or more prosodic domains, with those occurring at the beginning of more domains exhibiting more fortition (i.e., phrase-initial /s/ is also word-initial and syllable-initial, while syllable-initial /s/ is only syllable-initial). The third dimension of Table 11, *preferred prosodic positions,* shows that [s] prefers the most prosodically prominent/perceptually salient positions; [s^{θ}] prefers word-medial, onset position, which is slightly less salient; and [h] prefers syllable-final position, the least salient of all. Here we see that, while [s^{θ}] and [h] prefer similar surrounding phonological environments, their complementary distributions are determined by the prosodic positions in which they tend to occur.

Now that the theoretical stage has been set and the inclusion of [s^{θ}] in the lenition framework justified, the next chapter describes the methodology used to collect, analyze, and model data in this dissertation, including methods relevant to the present chapter.

Chapter 5. Methodology and data collection

5.1 Introduction

In order to investigate patterns of /s/ weakening in Salvadoran Spanish, I spent five weeks conducting fieldwork in El Salvador during August and September of 2015. This trip was made possible thanks to generous funding from the UCLA Ben and Rue Pine Research Travel Grant in addition to various Salvadoran contacts in Los Angeles who kindly connected me with friends and resources in their home country. All Salvadorans who participated in this project received monetary compensation for their time.

Now that Chapters 2 and 3 have situated the present study within the existing literature and theory and Chapter 4 has made a case for the inclusion of [s^{θ}] in this account of Salvadoran /s/ lenition, the present chapter's goal is to describe, in detail, the methodology used in this dissertation to collect and analyze the Salvadoran speech data. This chapter serves a variety of purposes; however, its primary goals are to provide the reader with a replicable procedure for collecting and analyzing naturalistic speech data within my proposed framework and to clarify and/or justify particular methodological decisions made along the way.

This chapter is divided into five sections. In Section 5.2, I discuss my participants, their demographics, and the recruitment and recording procedures in which they participated. In Section 5.3, I detail the process of segmenting, coding, and acoustically analyzing the speech data. Next, in Section 5.4, I describe the procedure for procuring lexical frequency data. In Section 5.5, I lay out the steps I took to convert my raw dataset into 14,400 OT-style tableaux and to procure constraint weights and output predictions via the Maxent Grammar Tool. Finally, Section 5.6

describes the procedure used to calculate adjustments to Faithfulness constraints based on different macro-social groups in order to incorporate social factors into the grammar.

5.2 Participants and procedure

Participants for this study were 72 speakers of Salvadoran Spanish residing in El Salvador at the time of data collection. While many participants reported elementary to intermediate knowledge of English, I consider this group of speakers as a whole to be monolingual speakers of the Salvadoran dialect. All participants were recruited via random and snowball sampling and were balanced for region, rural vs. urban origin, age, and gender, as illustrated in Table 12.

Table 12

Participant demographics

Region	Western (Santa Ana)								Central (San Salvador)								Eastern (San Miguel)							
Origin	rural				urban				rural				urban				rural				urban			
Age	18-40		41+		18-40		41+		18-40		41+		18-40		41+		18-40		41+		18-40		41+	
Gender	M	F	M	F	M	F	M	F	M	F	M	F	M	F	M	F	M	F	M	F	M	F	M	F
Total N=72	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3

Participants were recruited from three regions within El Salvador: the western region, the central region, and the eastern region. This division is consistent with how the regions are divided for statistical and geographical purposes by the national government ("Población y estadísticas demográficas," 2013) and reflects the hypothesis that the speech in the capital of San Salvador (pertaining to the central region) represents the national dialectal standard and should therefore be treated as discrete from other regions. Additionally, because of its close proximity to Guatemala, whose dialect generally retains /s/, I hypothesized that the western region might show evidence of dialect contact. Finally, Azcúnaga López (2010) identifies the eastern zone as a unique dialectal region based on the presence or absence of a number of phonological features. Confirming these findings, information obtained both in sociolinguistic interviews and in personal conversations with Salvadorans in the community revealed pervasive stereotypes associated with the speech of those from the eastern region.

I selected the largest city within the western, central, and eastern regions (Santa Ana, San Salvador, and San Miguel, respectively) from which to sample urban participants. Rural participants for a given region were sampled from rural areas within the department in which the city was located, also termed Santa Ana, San Salvador, and San Miguel. El Salvador is divided into fourteen geographical Departments, as illustrated in Figure 25.

Figure 25. El Salvador is divided into fourteen distinct departments; participants for the present study came from those of Santa Ana, San Salvador, and San Miguel (map from Wikipedia.com).

Given that the migration of rural Salvadorans to urban areas has increased in recent years in the wake of the civil war and intense gang violence, city-dwelling participants were only considered *urban* if they were born in the city in which they currently lived and had not resided elsewhere for more than six months. Prospective participants who did not meet these criteria were not recruited.

Age groups for the present study were established based on Eckert's (1997) emic approach, which relies on shared experiences such as social, political, and historical events to group speakers. Such an approach hinges on the idea that significant life experiences not only play a role in shaping identity but are also at the crux of social changes that are linked to linguistic innovation and change. This approach was deemed most appropriate given the effects of sociopolitical turmoil that began with the Salvadoran Civil War in the late 1970s and continue today in the form of widespread gang violence. The two age groups in this study—those who were older than 41 at the time of recruitment and those between 18 and 41 years of age—represent Salvadorans born before and after the war, respectively. As explained by many of my participants, crucial societal changes that have occurred as a result of the civil war and the violence that has followed, such as the mass migration of young males to North America and dramatic urbanization of the Salvadoran population, have created an important generational divide.

For reasons of feasibility, participants were not balanced for level of education. Instead, those who were determined to have met other demographic criteria were asked to disclose the highest level of education they had completed. Answers ranged from no primary education (n = 8) to a university degree (n = 9) and were subsequently divided into two meaningful groups according to literacy achievements and economic opportunities afforded by particular educational milestones.[24] The highest educational group included those with some secondary education or more (n = 39), and the lower educational group included those with some primary education or less (n = 33).

Speakers participated in a sociolinguistic interview lasting 45-60 minutes. The purpose of the sociolinguistic interview was, as always, twofold: I aimed to forge a more complete

[24] This information was elicited in sociolinguistic interviews and was further confirmed by local contacts in the community.

understanding of the social, political, economic, and cultural forces at work in El Salvador that might be affecting linguistic change and variation while simultaneously obtaining naturalistic speech samples from my participants. I began each interview with personal questions that would reveal demographic information about the speaker as well as build rapport. I would then transition to a series of questions about the speaker's immediate community, and then to broader issues about the state of the country in general, the civil war, the pervasiveness of gang violence and the role migration plays in modern Salvadoran society. I finished each interview with questions about language attitudes in an effort to ascertain what linguistic stereotypes speakers were aware of and which social groups were associated with each. Sample questions from each of these sociolinguistic interview sections appear in Table 13 below.

Table 13

Sample questions from sociolinguistic interviews

Topic: Work
¿A qué se dedica usted? ¿Le gusta? ¿Por qué (no)? ¿Cómo son sus compañero/as de trabajo? ¿Ha trabajado usted en otros oficios? ¿Cuáles le han gustan más y por qué?
What do you do for work? Do you like it? Why or why not? What are your workmates like? Have you worked in other professions? Which have you liked the most and why?
Topic: Community
¿Qué tipos de personas viven por esta zona? ¿Conoce bien a sus vecino/as? ¿A qué se dedican ellos? ¿Cómo ha cambiado este barrio durante su vida? ¿Hay algún lugar cercano donde la gente pasa sus ratos libres (un parque, una iglesia, etc.)?
What type of people live around here? Do you know your neighbors well? What do they do for work? How has this neighborhood changed over the course of your life? Is there somewhere nearby where people spend their free time (a park, a church, etc.)?
Topic: Social, political, and economic issues affecting El Salvador
¿Cuáles son los problemas más graves en El Salvador hoy en día? ¿Cómo influyen en la vida de usted? ¿Cómo ha cambiado el país en las últimas décadas? ¿Se van muchas personas para los Estados Unidos? ¿Por qué?
What are the most serious problems affecting El Salvador today? How do they affect your life? How has the country changed in recent decades? Do a lot of people leave to go to the U.S.? Why?
Topic: Language attitudes
¿Hablan diferente los habitantes de distintas regiones de El Salvador? ¿Hay algunos salvadoreños que hablan con acento? ¿Hablan diferente las personas con más/menos educación? ¿En qué sentido? ¿Hablan diferente los jóvenes y las personas mayores? ¿Las mujeres y los hombres? ¿En qué sentido?
Do people from different regions in El Salvador speak differently? Are there some Salvadorans who speak with an accent? Do people with more/less education speak different? How so? Do young people and older peoople speak differently? How about women and men? How so?

While I was careful to cover each of the above topics at least briefly in every interview, I was also committed to engaging with my participants in a meaningful way and therefore asked guided follow-up questions based on their responses as often as possible. All participants were recorded using an Olympus LS-14 Linear PCM recorder digitized at 44.1 kHz and a 16-bit sampling rate with an attached *Audiotechnica* ATR 3350 lapel microphone. All recordings took place in quiet environments at participants' homes, local universities, or community centers. The following section details how these recordings were segmented, coded, and acoustically analyzed.

5.3 Segmentation, coding, and acoustic analysis

Beginning 20 minutes[25] into each recording, the waveform and spectrogram of the first 200 occurrences of phonological /s/ (n = 14,400) were segmented in *Praat* (Boersma & Weenink, 2018). Tokens were excluded (i.e., passed over) in the following situations:

(a) immediately preceding or following a voiceless fricative (/s/, /x/, /t͡ʃ/, and /f/);

(b) in cases of disfluency;

(c) when there was background noise;

(d) when the participant and I spoke at the same time;

(e) when it was not clear from the context if the word in question had an underlying /s/ or not;

(f) when the word was a filler, such as *pues* ('well'), *este* ('um'), or *o sea* ('I mean'); and

(g) when the /s/ in question was the word-final sibilant in *quizás* ('perhaps'), as this word is used interchangeably with *quizá*.

For the purposes of this dissertation, the term *segmentation* refers to either a) placing boundaries around sections of the speech stream during which a phonological /s/ occurs, or b) indicating the absence of a segment at a particular moment where /s/ exists underlyingly but is not realized phonetically (i.e., a deleted token). As can been seen in Figure 26 below, tokens for which there was some phonetic realization of /s/ (be it [s], [z], [s$^{\theta}$], [h], or [ɦ]) were segmented on an interval tier so that acoustic measures such as duration could be easily extracted, while deleted

---

[25] While ten minutes is standard practice in sociolinguistics, challenges posed by a precarious fieldwork environment often required time during the first half of the interview to adjust and minimize ambient noise and interruptions.

tokens were marked on a point tier. For tokens segmented in the interval tier, I placed boundaries at the onset and offset of the fricative noise.

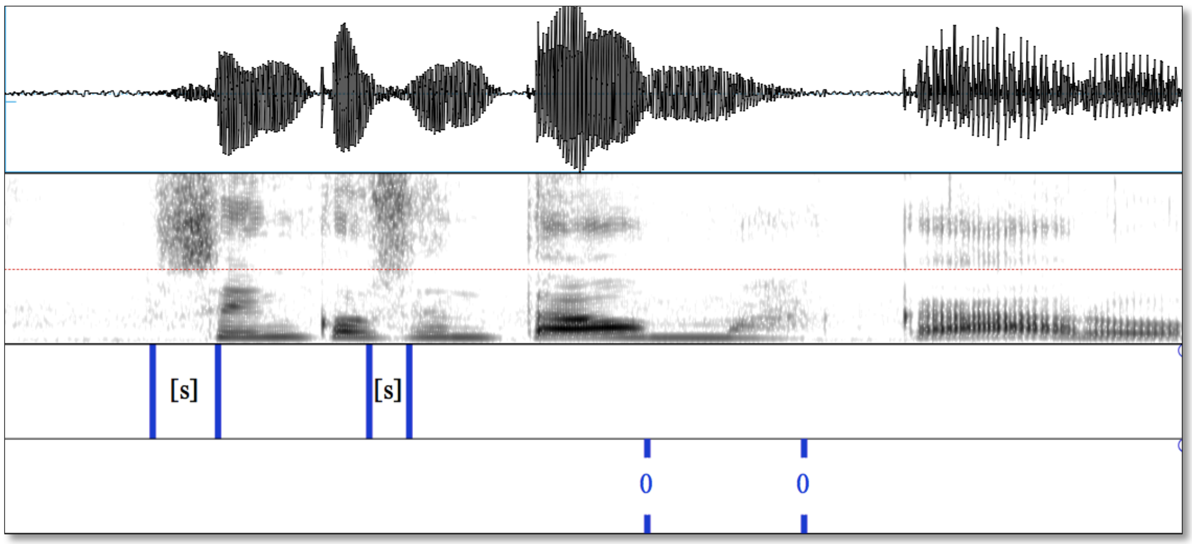

Figure 26. Spectrogram and waveform of the phrase cinco centavos más caro ('five cents more expensive') with the two occurrences of phonological /s/ in [s]inco and [s]entavos segmented on the interval tier and the two occurrences of /s/ in centavo[Ø] and má[Ø] marked on the point tier.

Given the variability of phonetic realizations of /s/, I was careful to be precise and methodical both in how I segmented fricatives using the waveform and spectrogram and in how I determined when a phonetic realization was entirely absent and should be coded as a deletion. Indeed, because speech sound segmentation is inherently difficult, choosing consistent measurement points is crucial (Ladefoged, 2003, p. 104). I therefore depended upon reliable acoustic properties in tandem with the following guidelines, adapted from Erker (2012, p. 49):

1. Fricative sounds have reliable and characteristic acoustic correlates that can be identified in spectrograms and waveforms.
2. Some or all of these correlates are present in non-deleted /s/ tokens.
3. /s/ 'deletion' can be defined as the absence of these correlates.

In order to precisely segment a given token of phonological /s/, I first determined the perceived allophone via a series of acoustic criteria useful for differentiating fricatives (see Johnson, 2012, pp. 152–162). Sibilants—that is, [s] and [z]—were identified by the presence of fricative noise at high frequencies and amplitudes, produced by turbulent airflow through a tight oral constriction near the front of the vocal tract. These sounds "display a primary spectral peak…around 4000 to 5000 Hz" (Jongman et al., 2000, p. 1253) as well as "random noise extending well beyond the upper limits of the spectrogram" (Ladefoged, 2003, p. 202). The voiced and voiceless sibilants were further differentiated from one another by the presence or absence of a voice bar as well as "the addition of vertical striations indicative of voicing" (Ladefoged, 2003, p. 202). Because voicing of /s/ is gradient (Campos-Astorkiza, 2014), a token was only coded as [z] if a voice bar was present throughout the duration of the fricative. Figure 27, Figure 28, Figure 29, and Figure 30 compare waveforms and spectrograms of [s] and [z] in order to illustrate these criteria and establish segmentation conventions. Note that sample spectrograms in this section appear in pairs—one token occurring in syllable onset position and the other in syllable coda position—for illustrative purposes.

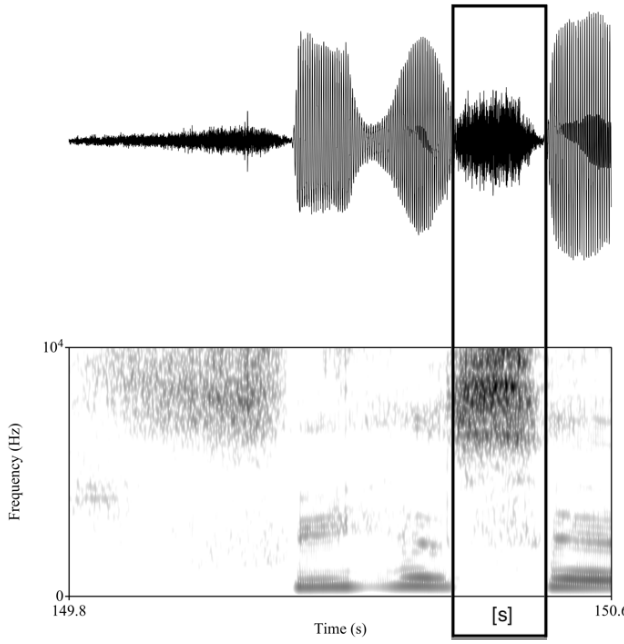

Figure 27. Sample spectrogram and waveform of onset [s] in *fijese* [ˈfi.xe.se] (‘look’).

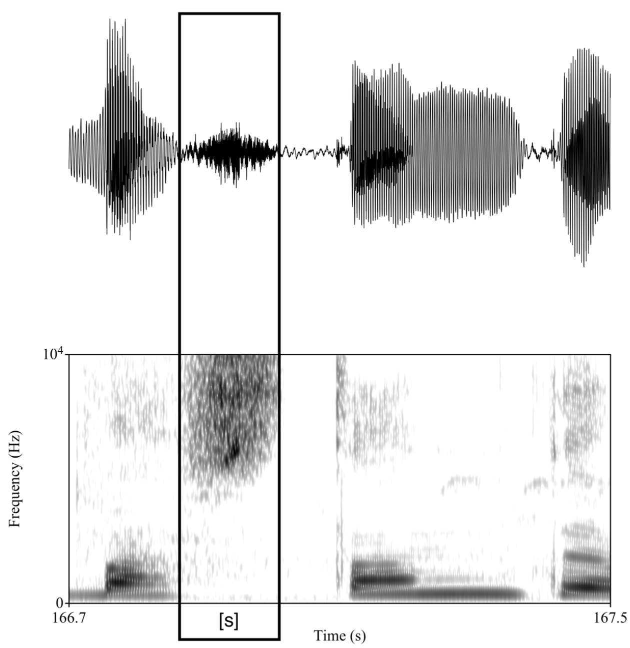

Figure 28. Sample spectrogram and waveform of coda [s] in *bastante* [basˈtan.te] (‘quite’).

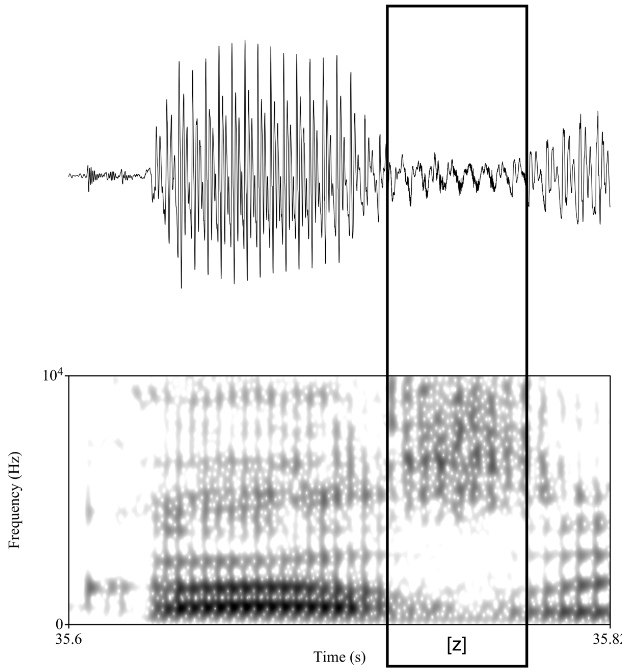

Figure 29. Sample spectrogram and waveform of onset [z] in *casa* [ˈka.sa] (‘house’).

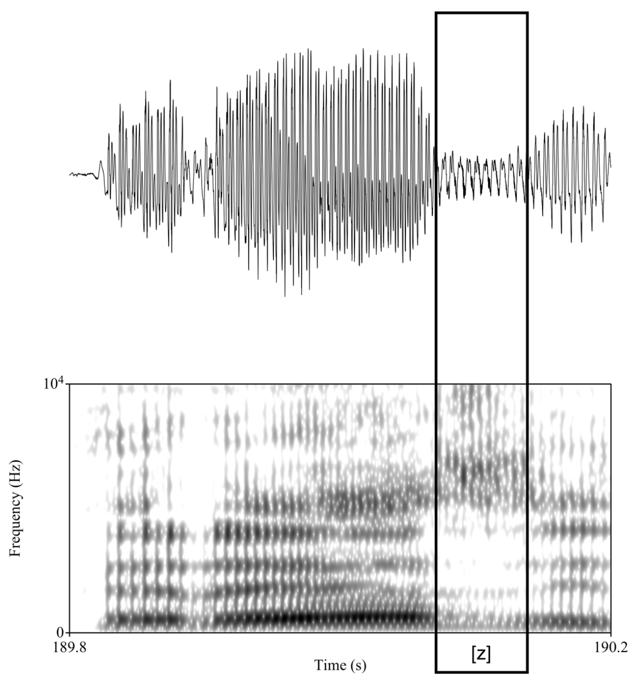

Figure 30. Sample spectrogram and waveform of coda [z] in *querés* [keˈɾez] (‘you want’).

Tokens of [s^{θ}] were identified and subsequently segmented based on the assertion that this sound results from some amount of gestural undershoot of /s/ (see Chapter 4). I therefore expected tokens of [s^{θ}] to have lower observed amplitude, due to the relaxation of the constriction between the tongue and the alveolar ridge as well as the slight separation of the upper and lower teeth created by the postdental position of the tongue. Similarly, I expected the frequencies of [s^{θ}] to decrease as compared to [s], as the absence of obstruction by the upper and lower teeth results in a longer front cavity and, subsequently, lower resonant frequencies (Johnson, 2012, p. 157). Figure 31 and Figure 32 show tokens of [s^{θ}] in onset and coda positions, respectively, characterized by fricative noise that is lower in both frequency and amplitude than its sibilant counterparts.

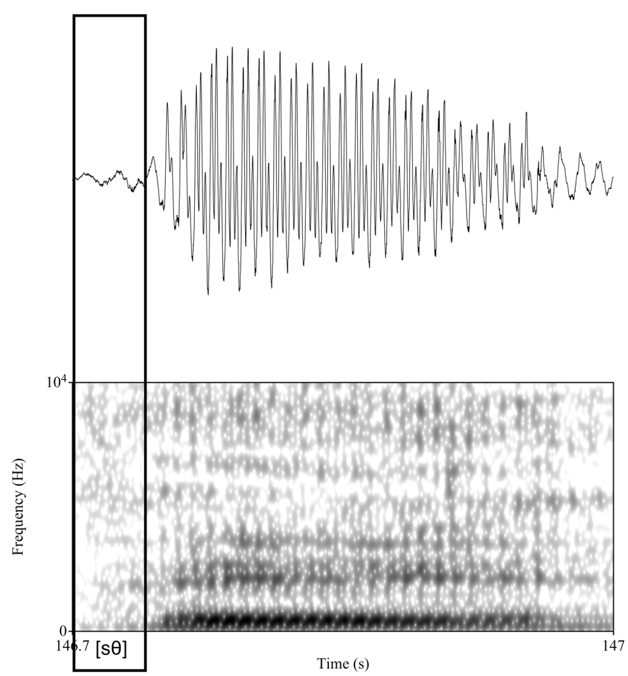

Figure 31. Sample spectrogram and waveform of onset [s^{θ}] in *sí* [s^{θ}i] (‘yes’).

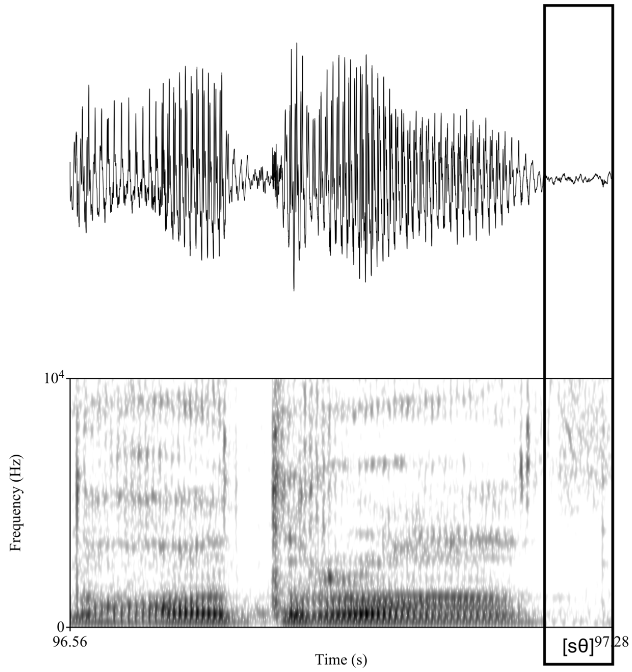

Figure 32. Sample spectrogram and waveform of coda [s^{θ}] in *nosotros* [noˈɦo.tɾos^{θ}] (‘we/us’).

Finally, tokens of the glottal fricative were distinguished from other fricative allophones by an absence of high-frequency noise and the presence of formants similar to those of neighboring vowels. In order to segment voiceless [h] and voiced [ɦ], I relied on aperiodicity in the waveform as well as the relative strength of formants visible in the spectrogram, both of which are key differentiators between glottal fricatives and vowels. Again, the presence of vertical striations and a voicing bar along the x-axis of the spectrogram were important in differentiating voiced and voiceless variants. These characteristics and criteria are illustrated in Figure 33, Figure 34, Figure 35, and Figure 36.

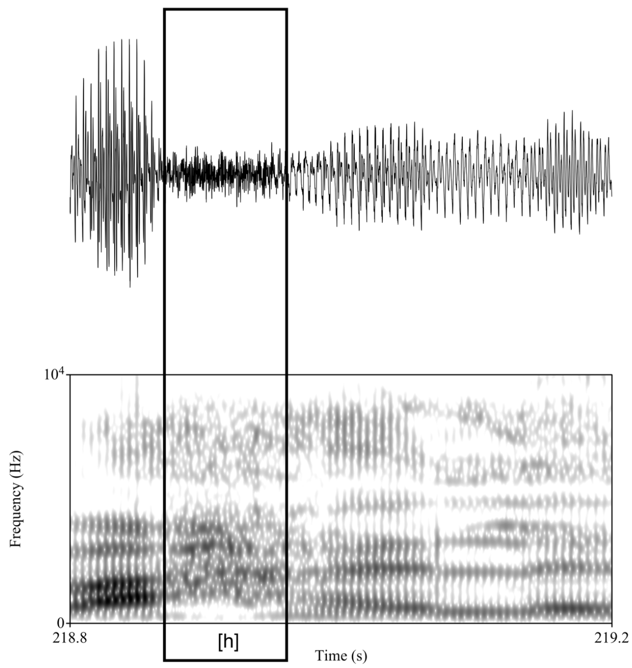

Figure 33. Sample spectrogram and waveform of onset [h] in *se* [he] (‘itself/themselves’).

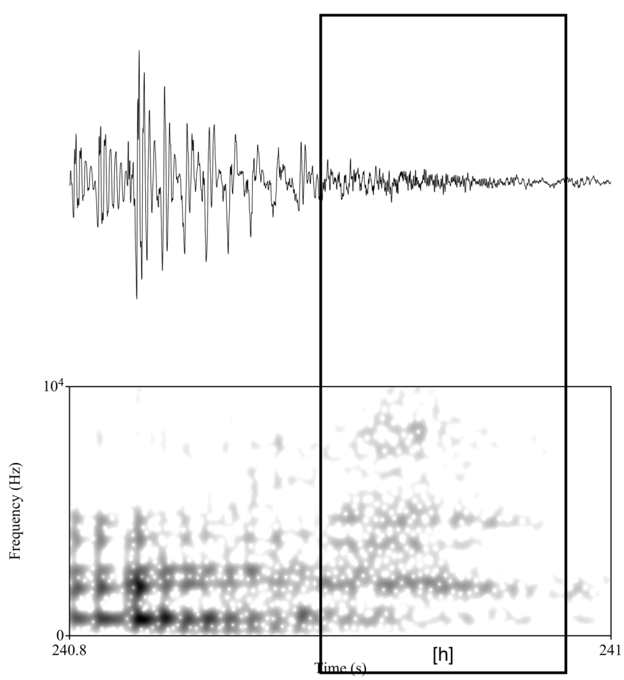

Figure 34. Sample spectrogram and waveform of coda [h] in *es* [eh] (‘is’).

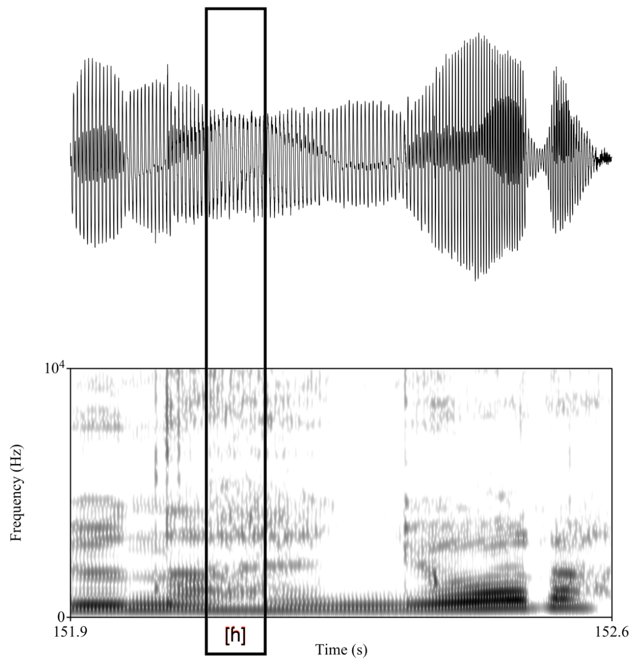

Figure 35. Sample spectrogram and waveform of onset [ɦ] in *señora* [ɦeˈɲo.ɾa] (‘ma’am’).

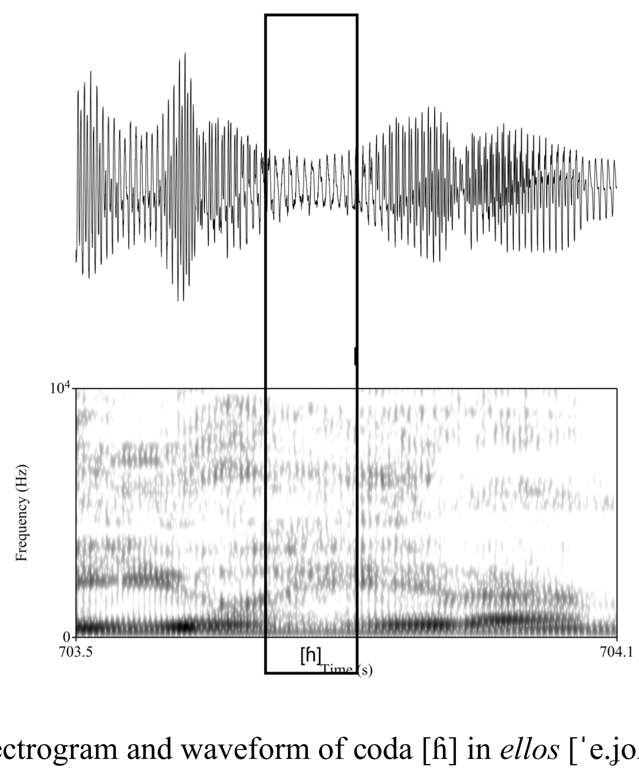

Figure 36. Sample spectrogram and waveform of coda [ɦ] in *ellos* [ˈe.ʝoɦ] (‘they/them’).

Finally, because determining the presence vs. absence of a consonant in a methodical way can be challenging, I followed guidelines established by Keating et al. (1994, p. 92):

> Sometimes a consonant is heard but…no clear interval that could be uniquely associated with the consonant is seen in the signal. If there is a brief local amplitude drop in the signal which could be assigned to the consonant, the consonant is transcribed. If there is no amplitude drop in the signal and/or no interval in which that consonant seems to predominate, then no consonant is transcribed. […] Occasionally a fricative, e.g. [s], is clearly heard, but only a gap is seen in the acoustic displays. In these cases the fricative is transcribed.

With these standards in mind, tokens for which all fricative correlates were absent and there was no amplitude drop in the signal were coded and marked as deletions. It should be noted that, while Lipski (1984, 1985) does not find any cases of /s/ deletion in syllable-initial position in his Salvadoran data, onset deletion was present (albeit very rare) in mine (n = 65 tokens). Figure 37 and Figure 38 show sample deletions of onset /s/ and coda /s/, respectively.

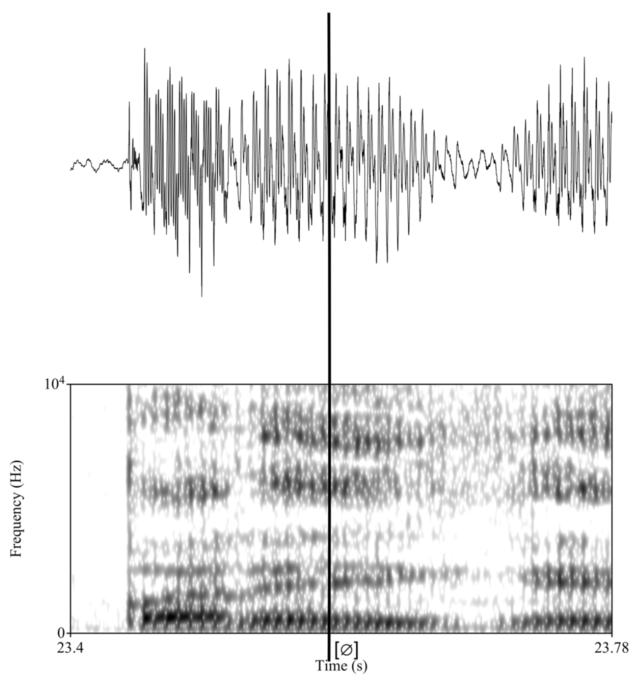

Figure 37. Sample spectrogram and waveform of onset deletion in *parece* [paˈɾeː] (‘it seems’).

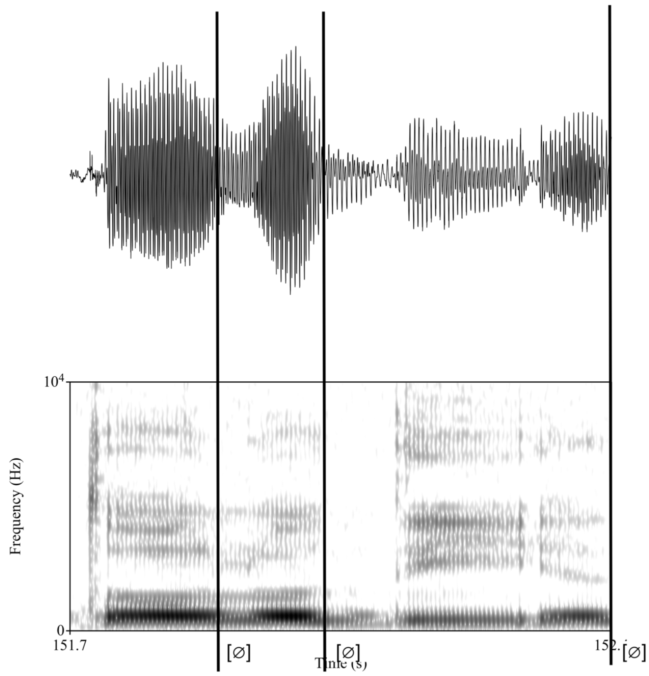

Figure 38. Sample spectrogram and waveform of three sequential coda deletions in *todos los días* [to.loˈdi.ja] (‘every day’).

In Figure 37, the reader can see that the word *parece* is pronounced with a long [e:], as the intervening /s/ is deleted entirely. While there is a very slight drop in amplitude halfway through this long vowel, indicated by a line in through the waveform and spectrogram, there is no clear interval during which a consonant might be segmented or even reliably identified.

In addition to transcribed allophone, each token of /s/ was coded for the following social and linguistic factors: region, rural/urban origin, age, sex, education, prosodic position, preceding segment, following segment,[26] syllable stress, and stress of the following syllable. To facilitate the identification of phrase boundaries, i.e., the beginning/end of an intermediate or intonational phrase, I used the F0 tracker function in *Praat* to identify boundary tones in conjunction with

[26] For the factors of *preceding* and *following segments*, glides resulting from diphthongization, such as the /j/ in *canción* /kansjon/ (‘song’) and the /w/ in *sueño* /sweɲo/ (‘dream’), were coded as high vowels.

measurements of pause length[27] and evidence of final lengthening of the syllable, even if /s/ was absent. Table 14 below shows the independent variables coded for in my analysis and their respective levels.

[27] Yang (2007) shows that in conversational speech cross-linguistically, pause duration is significantly correlated with boundary status, with major boundaries (such as those marking intermediate and intonational phrases) averaging 460 ms and non-boundary pauses averaging 280 ms.

Table 14

Independent variables coded for analysis

Factor	**Levels**
Region	Santa Ana
	San Salvador
	San Miguel
Origin	Urban
	Rural
Age	18-41
	41+
Sex	Male
	Female
Education	Level 1
	Level 2
Prosodic position	Phrase-initial
	Word-initial
	Syllable-initial
	Syllable-final
	Word-final
	Phrase-final
Preceding segment	vowel $_{[+high]}$
	vowel $_{[-high]}$
	consonant $_{[+coronal, +voice]}$
	consonant $_{[+coronal, -voice]}$
	consonant $_{[-coronal, +voice]}$
	consonant $_{[-coronal, -voice]}$
	pause
Following segment	vowel $_{[+high]}$
	vowel $_{[-high]}$
	consonant $_{[+coronal, +voice]}$
	consonant $_{[+coronal, -voice]}$
	consonant $_{[-coronal, +voice]}$
	consonant $_{[-coronal, -voice]}$
	pause
Syllable stress	tonic
	atonic
Following syllable stress	tonic
	atonic
	pause (N/A)

In addition to these independent variables, the phonological word in which /s/ occurred, the participant's ID number, and a token number were also provided for each token.

Once all 14,400 tokens had been segmented and coded, I ran a Hann low pass-band filter in *Praat* that eliminated frequencies below 750 Hz for all non-zero tokens. Following File-Muriel and Earl Brown (2010), I did this prior to extracting acoustic measures in an effort to capture the noise component of the fricative while excluding glottal pulsing. Duration and COG measurements[28] for non-zero tokens were then extracted using a script developed by Lars Hinrichs (Department of English, UT Austin) that measures COG at three points: one quarter, one half, and three quarters of the way through the segment. Following Erker (2010), I assigned deleted tokens a duration of '0' and a COG of 'undefined.'

In addition to the independent variables listed in Table 14, I investigated the effect of lexical frequency on rates of lenition. The following section describes the process for obtaining lexical frequency data and incorporating it into my analysis.

5.4 Lexical frequency data

Lexical frequency data for this dissertation was obtained from Brigham Young University's *Corpus del Español (CdE)*, which is available at https://www.corpusdelespanol.org/web-dial/ (Davies, 2016). This 2-billion-word corpus, funded by the U.S. National Endowment for the Humanities, uses data from approximately two million web pages from 21 Spanish-speaking countries. The *CdE* was deemed more appropriate for the present study than the Real Academia's *CORPES* ("Corpus del Español del Siglo XXI

[28] COG is a measure of the average frequency of fricative energy and can be calculated by the equation $COG = \sum fI / \sum I$, where I is amplitude (in dB) and f is frequency (in Hz) (Erker, 2010:13). See Section 2.4 for a detailed description of this measure.

(CORPES)," n.d.) for a number of reasons, including size and dialectal representation. The *CdE* is not only close to 12 times larger than *CORPES*, but also reflects a more faithful representation of the population balance between Latin America and Spain: while only 10% of the world's Spanish-speaking population is from Spain, approximately 35% of *CORPES* data comes from the peninsular dialect (as compared to only 22% of *CdE*'s). Furthermore, El Salvador is fairly well-represented in the *CdE*, contributing a total of 39,147,035 words to the corpus.

While projects such as Erker's (2012) account of /s/ weakening in New York City look at both host-word and lemma frequencies, this dissertation focuses solely on the former due to the large number of tokens analyzed as compared to studies of a similar nature.[29] Distinguishing between two host-words with the same lemma (e.g., *hace* and *hiciera,* both corresponding to the lemma *hacer*) is important, as frequency data can be extremely disparate in these cases. I therefore conducted a lexical frequency search for each unique word in the dataset (n = 1975) and requested additional results specific to the Salvadoran section of the corpus in the form of occurrences per million. Figure 39 shows sample search results.

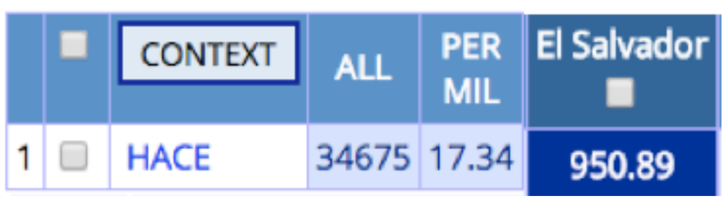

		CONTEXT	ALL	PER MIL	El Salvador
1		HACE	34675	17.34	950.89

Figure 39. Sample search results for the host-word *hace* ('he/she/it does/makes') from BYU's Corpus del Español (CdE).

[29] For instance, in recent dissertations on Spanish /s/ weakening, Erker (2012) and Chappell (2013) examine 4,800 and 7,133 tokens of /s/, respectively.

In Figure 39, the column "ALL" indicates the total number of occurrences of this word in the entire corpus, the column "PER MIL" indicates the occurrences of this word per million words in the entire corpus, and the column "El Salvador" represents the occurrences of this word per million words in the Salvadoran section of the corpus only.[30] This third number, 950.89, was entered into the dataset as the raw lexical frequency of the word *hace*. Lexical frequency data in the form of occurrences per million words in El Salvador ranged from 0.00 words per million (lowest lexical frequency) to 17,401.74 words per million (highest lexical frequency). These raw values were then converted to log frequencies, with raw values of 0.00 rounded up to 0.001. The distribution of log frequencies can be seen in Figure 40.

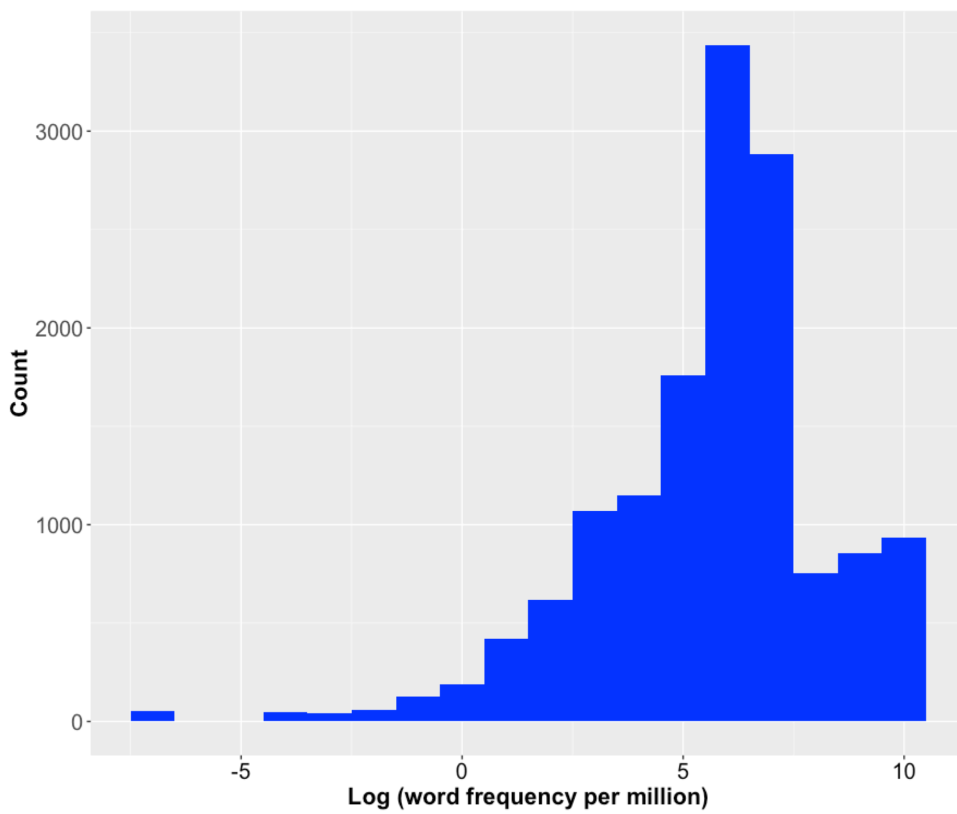

Figure 40. Log frequency (based on words per million) distributions for the Salvadoran data.

[30] Because there is significantly less lexical variation in each dialectal faction of the corpus, there were often large discrepancies between the "ALL" and "El Salvador" columns.

Here, we see that the majority of words in the dataset are most densely clustered between log frequencies of 5 and 7.5 and range from approximately -7 to 10. Consistent with methodological norms in corpus linguistics, log frequency was determined to be the most accurate measure because differences in raw frequencies between high-frequency words is much larger than differences between low-frequency words, yet the size of these differences is not necessarily meaningful with respect to word recognizability or confusability. For example, the difference in raw frequency between the highest-frequency word in the dataset, the masculine plural article *los*, and the second-highest-frequency word, the reflexive pronoun *se*, was 2,453.49 words per million. Using raw frequencies would suggest that the difference in recognizability between these words was extremely large, which is a gross distortion of the data. The difference in log frequencies of these two words is 0.15, which is a much more valuable measure of the minute frequency difference between these two extremely high-frequency words.

This section has described the process for procuring and transforming lexical frequency data. The importance of this factor in conditioning patterns of /s/ lenition will be explored in Section 6.2. The following section describes the procedure for formatting the data in OT-style tableaux and training the maxent grammar.

5.5 The maxent grammar

In order to address the most fundamental research questions of this dissertation, I built a Maximum Entropy model of Salvadoran /s/ lenition trained on my 14,400 tokens of naturalistic speech data using the Maxent Grammar Tool (Hayes et al., 2009). In this section, I walk the reader through the procedure for creating the input file that was used to train the grammar as well as the process for calculating constraint weights and probabilities for known and novel output forms.

5.5.1 The Maxent Grammar Tool

The Maxent Grammar Tool is a resource for linguists interested in modeling language learning by developing a grammar trained on real speech data. The linguist must develop the grammatical constraints herself and provide the program with the following information: underlying representations ("inputs") for unique tokens and rival surface representations (or, "outputs") for each, indications of winning outputs and their frequencies, a list of pertinent constraints, and the number of violations incurred by each candidate for each constraint. Figure 41, taken from the Maxent Grammar Tool Manual (Hayes, 2009), shows a sample input file.

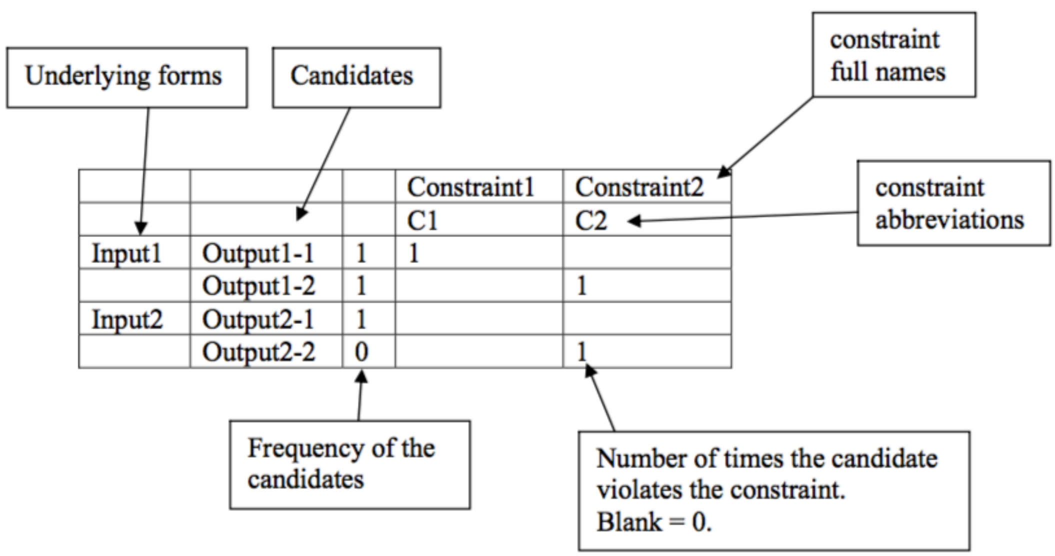

Figure 41. Sample input file for the Maxent Grammar Tool, from Hayes (2009).

In order to convert my dataset from a standard rows and columns format into an input file like that required by the Grammar Tool, I first expanded each underlying form into a six-row tableau and developed a complete set of constraints (n = 137),[31] which I then placed along the top row of the input file. Next, I wrote a series of "if" statements to translate information in the dataset into constraint violations. For example, the constraint LAZY(vcl_strid_fric,$V_{[+high]}$__$V_{[+high]}$) incurs one violation for every instance of [s] between two high vowels, so I wrote an "if" statement indicating that one violation should be tabulated for every co-occurrence of a candidate containing allophone [s], preceding segment [i] or [u], and following segment [i] or [u]. The software then uses these violations in tandem with observed values of each candidate to calculate the best constraint weights.

Once an input file has been uploaded to the Maxent Grammar Tool, the program uses the tableaux to a) produce an optimum grammar to match the training data by assigning weights to the constraints it's given, and b) make quantitative predictions for known and novel forms based on these constraint weights. Figure 42, again borrowed from the Grammar Tool Manual, shows a sample output file, which Hayes has color-coded for clarity.

---

[31] A complete list of constraints used to develop this grammar can be found in Appendix A.

Input: Input1		C1	C2
Output1-1	1	1	0
Output1-2	1	0	1
Input: Input2		C1	C2
Output2-1	1	0	0
Output2-2	0	0	1

|weights| after optimization:			
Constraint1 (mu=0.0, sigma^2=100000.0)	8.041368461		
Constraint2 (mu=0.0, sigma^2=100000.0)	8.041690116		

Input:	Candidate:	Observed:	Predicted:
Input1	Output1-1	1	0.500080414
Input1	Output1-2	1	0.499919586
Input2	Output2-1	1	0.999678339
Input2	Output2-2	0	3.22E-04

Figure 42. Sample output file for the MaxEnt Grammar Tool, from Hayes (2009).

As the reader can see, the output file contains three distinct types of information. In red, we see a copy of the original tableaux; in blue, we see the weights assigned to each constraint; and in green, we see predicted probabilities assigned to candidates. This last section is particularly important: by comparing the “observed” and “predicted” columns of this last section of output, we can not only get a sense for how well the grammar fits the training data, but we can additionally obtain predicted probabilities for novel inputs that the grammar has never seen (new test forms may be included in the input file with frequencies of ‘0’ so that they are not used as training data and don’t affect the weights learned). This is akin to a real-life language learner who uses her implicit knowledge of a grammar to produce outputs for novel inputs. The following sections describe how the constraint weights and predicted probabilities reported by the Maxent Grammar Tool are determined.

5.5.2 Calculating constraint weights

The job of the maxent learning algorithm is to maximize the *predicted probability of the observed data* minus a *penalty for non-default weights* by adjusting constraint weights; this mathematical expression, called the objective function, appears at the end of this section. The Maxent Grammar Tool uses gradient ascent to calculate these optimum constraint weights, a process comparable to that used to determine weights in a logistic regression. Every constraint starts with the same initial weight—an arbitrary starting place—and the system calculates the probability of the data at this weight. The system then calculates the slope of the function at that stage, uses that information to move closer to the peak of the function, adjusts the weights, and recalibrates. The learner continues this process until it arrives at the peak of the function, or global optimum, where the probability of the data has been maximized. For a detailed and fairly straightforward discussion of gradient methods and their associated calculations, see Shewchuk (1994).

In addition to the probability that the grammar assigns to the observed data, the penalty assigned for non-default weights is a crucial component of the objective function. The purpose of this component, which penalizes weights that are either too large or are too different from the default value for a given weight, is to avoid overfitting the data (i.e., a model that fits training data too well might make worse predictions about new data) by ensuring that large and/or deviant constraints make an important contribution to explaining variation in the data. In the Maxent Grammar Tool, we can give each constraint its own default weight, μ. The amount a given constraint is allowed to deviate from that default weight is then specified by σ, with small values resulting in severe penalties for deviations from the specified μ and larger values enacting weaker penalties.

Keeping these components in mind, the mathematical expression of the objective function appears below. The first part of the equation denotes the log probability that the grammar assigned to the observed data, where x_i is a given set of data. The second part of the equation is a measure of constraint weight penalty, where w_j is the weight of constraint c_j, μ_j is the default weight for that constraint, and σ_j is the willingness of that constraint to deviate from μ_j.

$$\sum_{i=1}^{N} \ln P(x_i) - \sum_{j}^{M} \frac{(w_j - \mu_j)^2}{2\sigma^2}$$

In this expression, $P(x_i)$ is calculated by taking the sum of a given candidate's weighted constraint violations, making that sum negative and exponentiation it, and then dividing the resulting number by the negative, exponentiated sum of weighted constraint violations for all candidates in the tableau. In other words, $e^{-\text{SumWeightedConstraintViolationsCandidatej}} / e^{-\text{SumOfWeightedConstraintViolationsAllCandidates}}$. The sum of weighted constraint violations, referred to as a candidate's *harmony*, is fundamental to stochastic grammars and is revisited in the following section.

5.5.3 Calculating probabilities

Through the process of gradient ascent described above, every constraint in the maxent grammar is assigned a weight in the form of a non-negative real number, which serves to indicate its relative importance: constraints with higher weights are more influential in the grammar, and therefore outputs that violate high-ranking constraints are less probable. Once constraint weights have been determined by the Grammar Tool, it performs a series of steps to calculate the predicted probabilities of output candidates. These steps, which I have adapted from Hayes and Wilson (2008, p. 6) are as follows:

For a given output form *x*…

(a) Calculate the **harmony** of *x*, denoted here as h(x), which is the weighted sum of *x*'s total constraint violations:

$$h(x) = \sum_{i=1}^{N} w_i C_i(x)$$

where:
w_i is the weight of *i*th constraint
$C_i(x)$ is the number of times the *i*th constraint is violated by *x*
$\sum_{i=1}^{N}$ indicates that values should be summed over all constraints

(b) Calculate the **maxent value** of *x*, denoted here as P*(x)

$$P^*(x) = \exp(-h(x))$$

In other words, *e* (the base of the natural logarithm) is raised to the negative harmony of *x.*

(c) Calculate the **probability** of *x,* denoted here as P(x), by determining its share in the total maxent values of all output candidates, a quantity denoted as Z:

$$P(x) = P^*(x) / Z \qquad \text{where } Z = \sum_{y \in \Omega} P^*(y)$$

To summarize steps (a)-(c) above, a candidate's harmony is first calculated by multiplying each constraint's weight by the number of violations incurred by the candidate for that given constraint and summing those numbers. Next, the harmony is exponentiated ($e^{-harmony}$) and the share of total harmony (candidate's $e^{-harmony}$ divided by total $e^{-harmony}$) is calculated to find the candidate's probability. Crucially, because maxent maximizes the probability of distribution of a given variable, a candidate can never have a probability of 'zero,' and will only be assigned a probability of 'one' if it is the only candidate in the tableau.

In order to exemplify this process, I encourage the reader to examine Table 15, which shows sample harmonies, maxent values, and probabilities for six possible representations of the word *blusita* /blusita/ ('little blouse') based on imaginary weights assigned to two competing constraints: LAZY (vcl_strid_fric, $V_{[+high]}$__$V_{[+high]}$) ('Assign one violation for every [s] that occurs between two high vowels') and PRES(strid)/onset, tonic ('Preserve an input [+strident] specification in syllable-initial position in a tonic syllable').

Table 15

Scores and maxent values for six possible output forms of blusita */blusita/*

Output	LAZY (vcl_strid_fric, $V_{[+high]}$__$V_{[+high]}$) (*w*=2.0)	PRES(strid)/ onset, tonic (*w*=3.0)	*Harmony* (h(x))	*Maxent value* (P*(x))	*Probability* (P(x))
blu[s]ita	2.0 · 1 = 2.0	3.0 · 0 = 0.0	2.0 + 0.0 = 2.0	exp(-2.0) ≈ 0.14	0.14/ 1.34 ≈ 0.10
blu[z]ita	2.0 · 0 = 0.0	3.0 · 0 = 0.0	0.0 + 0.0 = 0.0	exp(-0.0) = 1.00	1.00/ 1.34 ≈ 0.75
blu[s^{θ}]ita	2.0 · 0 = 0.0	3.0 · 1 = 3.0	0.0 + 3.0 = 3.0	exp(-3.0) ≈ 0.05	0.05/ 1.34 ≈ 0.04
blu[h]ita	2.0 · 0 = 0.0	3.0 · 1 = 3.0	0.0 + 3.0 = 3.0	exp(-3.0) ≈ 0.05	0.05/ 1.34 ≈ 0.04
blu[ɦ]ita	2.0 · 0 = 0.0	3.0 · 1 = 3.0	0.0 + 3.0 = 3.0	exp(-3.0) ≈ 0.05	0.05/ 1.34 ≈ 0.04
blu[Ø]ita	2.0 · 0 = 0.0	3.0 · 1 = 3.0	0.0 + 3.0 = 3.0	exp(-3.0) ≈ 0.05	0.05/ 1.34 ≈ 0.04

In Table 15, we see that the Markedness constraint has been assigned a weight of 2.0 while the Faithfulness constraint has been assigned a weight of 3.0, indicating that the latter is more important than the former in this imaginary two-constraint grammar. These constraint weights in tandem with the number of violations incurred by a candidate are used to calculate the candidate's harmony, which is then used to compute its maxent value. Finally, a candidate's share of maxent values, i.e., its probability, is calculated by dividing its maxent value by the total sum of maxent values (in this case, 0.14 + 1.00 + 0.05 + 0.05 + 0.05 + 0.05, or 1.34). We see that blu[z]ita is by far the most probable outcome form in this imaginary grammar at 75%, which is logical given that it violates neither the Markedness constraint nor the Faithfulness constraint.

This section has provided a step-by-step account of the procedure used in this dissertation to build and train the maxent grammar, from preparing the input file to calculating probabilities for output forms using constraint weights. Next, in the final section of this chapter, I describe the procedure used to calculate adjustments to faithfulness constraints in order to incorporate social factors into the grammar.

5.6 Incorporating social factors

In order to incorporate social factors into the grammar, I implemented a procedure loosely adapted from that detailed in Coetzee's (2016) Noisy HG account of nasal place assimilation (see Section 3.8 for a detailed discussion of his methodology). Assuming that the effort cost associated with producing a given gesture or set of gestures in a given context does not change for different speakers, I follow Coetzee and scaled the importance of Faithfulness up or down according to social factors. Within this framework, differences in lenition between social groups depends on how much a given group prioritizes Faithfulness, i.e., the preservation of perceptual cues between the input and output, as compared to the Salvadoran speech community as a whole.

In order to do so, I developed independent maxent models for each individual social group (n = 11 sub-models) and compared the relative importance of Faithfulness in these sub-models with that of the base grammar. Relative importance of Faithfulness overall was calculated by taking the difference between average Faithfulness and Markedness constraint weights for each model, and the relative importance of specific Faithfulness constraints was calculated by taking the difference between the Faithfulness and Markedness constraint weights affecting a given aspect of Faithfulness. For example, the relative importance of preserving the input [+segmental] specification for urban speakers was calculated taking average constraint weights for PRES(seg)$_{UrbanModel}$ and subtracting average constraint weights for LAZY(any_segment)$_{UrbanModel.}$ Then, a scaling factor was calculated by subtracting the relative importance of preserving [+segmental] for urban speakers from its relative importance in the Salvadoran speech community as a whole. These methods will be further elucidated when social factor results are presented in Section 6.5.

To my knowledge, this is the first Optimality Theoretic account that incorporates the social characteristics of its speakers without building them into the grammar itself (Section 3.8 for a discussion of whether grammatical and non-grammatical factors should be treat the same or differently in the construction of a grammar). That is, the procedure I have described in this chapter yields a grammar-dominant model that serves as a foundation for the greater speech community while allowing for adjustments in constraint ranking by individual macrosocial groups.

5.7 Conclusion

Thus far, this dissertation has reviewed the literature on Spanish /s/ lenition, the phonetic and phonological theories that inform this process, and the methods employed to answer the research questions posed in Chapter 1. Specifically, the present chapter has described the methodology used to collect and analyze speech data, define linguistic and social variables, train the maxent grammar using real speech data, and incorporate social characteristics into the base model. The following chapter presents and explores the grammar developed by the Maxent

Grammar Tool as well as the scaled grammars that have been fitted to the data of individual social groups, using graphical representations of the data to illuminate the various models presented.

Chapter 6. Results: A maximum entropy model of Salvadoran /s/ lenition

6.1 Introduction

This chapter presents the maxent grammar for Salvadoran /s/ lenition and explores the results through graphical representations, raw data, and supplementary statistical analyses. All statistical analyses were conducted in R and all graphs that appear in this chapter were created using the ggplot2 package (Wickham, 2009) unless otherwise noted. Before presenting the base grammar, however, this chapter begins with an analysis of lexical frequency effects in the Salvadoran data to determine whether this factor merits inclusion in the final model. Then, in Section 6.3, I describe the process for determining the best value for the parameter σ to ensure that the model neither over- nor under-fits the observed data and is able to make reliable predictions about future data.

In Section 6.4, I describe the results of the base grammar developed by the maxent learning algorithm. First, in Section 6.4.1, I explore Markedness in the grammar context-by-context, first focusing on intervocalic contexts, then preconsonantal contexts, then C__V contexts, and finally, pause-adjacent contexts. Next, Section 6.4.2 examines Faithfulness in the grammar, this time by constraint type. In this section, I begin by presenting constraint weights for the twelve PRESRVE(strident, voice) constraints, and then examine additional penalties incurred for violations of PRESERVE(strident), PRESERVE(voice), PRESERVE(coronal), and PRESERVE(segmental). Next, in Section 6.4.3, I examine the interaction between Markedness and Faithfulness in the grammar via a series of tableaux that illustrate important features of configurations favoring oral constriction, configurations favoring debuccalization or deletion, and special cases that reveal particularly interesting facets of the grammar. Finally, in Section 6.4.4, I compare predicted and observed probabilities by allophone, context, and prosodic position in order to assess the grammar's strengths and weaknesses.

Once the features of the base grammar have been established and thoroughly examined, Section 6.5 looks at changes to the relative importance of Faithfulness for various groups of Salvadoran speakers. Section 6.5.1 looks at these changes with respect to region and makes important connections between constraint weights and patterns observed in the raw data. Tableaux are also presented to illustrate the implications of changes in the relative importance of Faithfulness on predictions made by the grammar for each regional group. Sections 6.5.2, 6.5.3, 6.5.4, and 6.5.5 follow a similar trajectory for the social factors of origin, age, gender, and level of education, respectively. Finally, Section 6.5.6 examines significant interactions effects between social groups on allophonic variation via multinomial logistic regression.

6.2 Lexical frequency effects

The effect of lexical frequency on /s/ weakening was difficult to hypothesize given the complexity of the lenition paradigm found in Salvadoran Spanish. That is, while some tripartite and instrumental accounts of /s/ weakening have identified lexical frequency as an important predictor, I was unsure how this effect would manifest itself in the Salvadoran data, if at all. In order to ascertain whether this factor plays an important role in determining the observed patterns in the data, I first examined the relationship between log frequency and allophone of /s/. This relationship can be visualized in Figure 43.

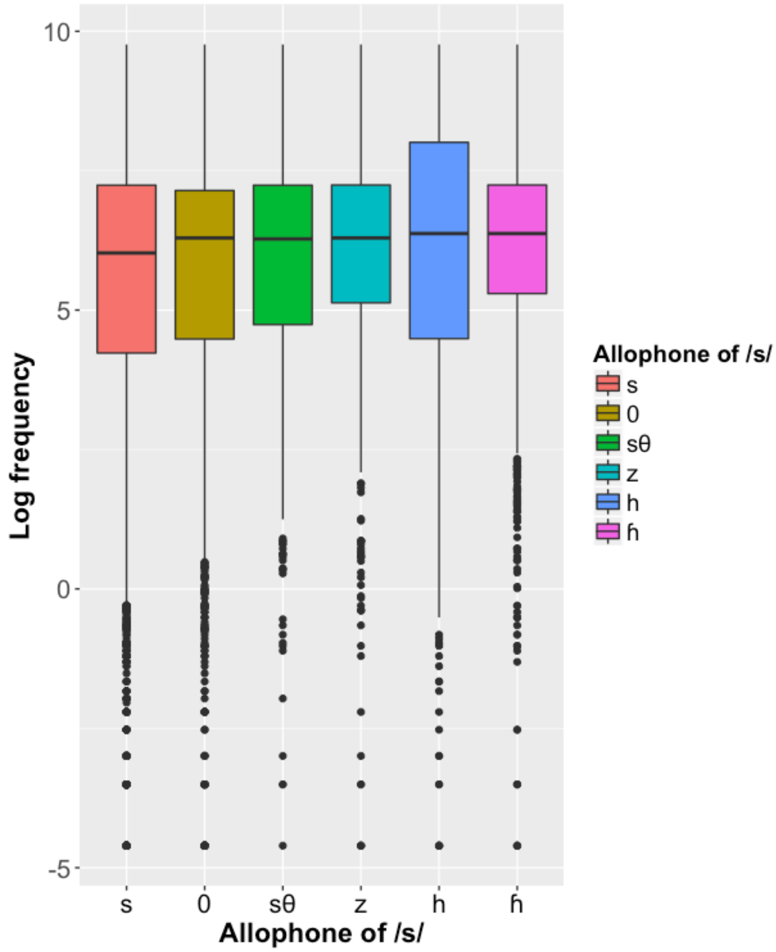

Figure 43. Allophone of /s/ by log lexical frequency, ordered from the lowest-frequency grouping to the highest-frequency grouping.

As the reader can see, the relationship between lexical frequency and allophone of /s/ is not clear-cut. On one hand, we see that [s] occurs in the lowest frequency words on average, supporting the hypothesis that lexical items with fewer lenited exemplars are less prone to lenition. However, we see that deletion is also prevalent in words of low frequencies, calling into question this same hypothesis. While a different ordering of these allophones might suggest an effect of lexical frequency on either the retention vs. loss of sibilance or the presence vs. absence of a segment, the observation that words containing [s] and [Ø] are, in fact, the most similar with respect to lexical frequency indicates that either this factor does not condition Salvadoran /s/ lenition, or we must attempt a more nuanced understanding of the data.

Acknowledging that lexical recall begins word-initially and that lexical confusion is, therefore, most at risk in this position (Keating, 2006, p. 180), I also examined differences in only this prosodic position.[32] Figure 44 shows lexical frequency by allophone for just word-initial prosodic positions, revealing clearer differences between words containing [s] and [∅] as well as those containing other, intermediate allophones.

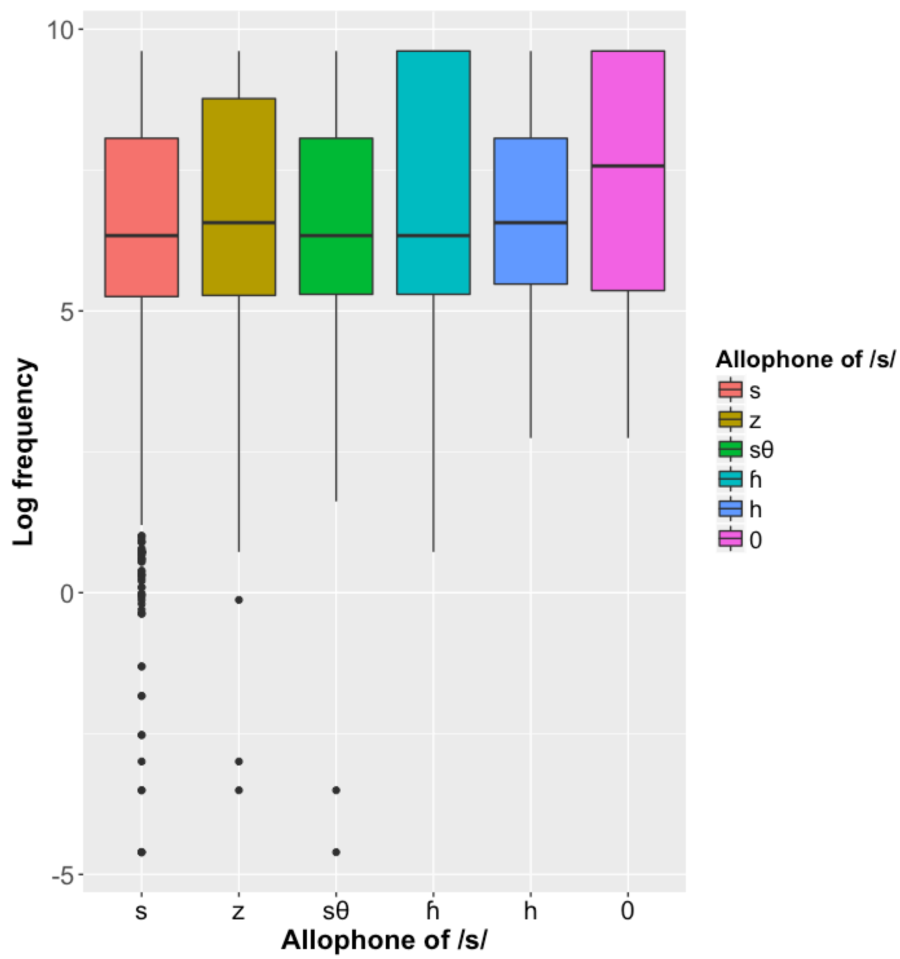

Figure 44. Allophone of /s/ by log word frequency for word-initial positions.

[32] For these purposes, both word-initial (phrase-medial) and phrase-initial tokens of /s/ were examined.

However, in order to ensure that these observed differences were not simply the artifact of sample size (allophones [h], [ɦ], and [∅] were much less common than [s], [z], and [s^{θ}] in this prosodic position), I ran a linear regression predicting log frequency based on allophone of /s/. The results of this regression, which can be seen in Table 16, reveal that no allophones are significantly different from [s] with respect to the dependent variable.

Table 16

Results of a linear regression comparing log frequencies for allophones of /s/

Allophone	**Estimate**	**Standard Error**	***t* value**	***p* value**
Intercept ([s])	6.42	0.05	140.78	<0.001***
[z]	0.19	0.22	0.90	0.37
[s^{θ}]	0.22	0.19	1.16	0.25
[ɦ]	0.32	0.19	1.65	0.10
[h]	0.48	0.38	1.28	0.20
[∅]	0.57	0.57	1.00	0.32

Given these results, lexical frequency was not used as a factor in the present analysis. The next section describes the procedure for determining the best parameters for the 137-constraint grammar containing 77 Markedness and 60 Faithfulness constraints.

6.3 Cross-validation for best fit

Before developing the complete base-grammar for Salvadoran /s/ lenition, I executed a smoothing procedure in order to avoid overfitting the final model to the training data. As I discuss in Section 5.5.1, a model that is too good of a fit for the data on which it is trained may do a worse job at making predictions about future, novel data. In the Maxent Grammar Tool, overfitting can be mitigated by making adjustments to the default values of μ and σ which, as described in the previous chapter, indicate default weights for a given constraint and a measure of that constraint's willingness to change, respectively. That is, making various adjustments to values for these two parameters tells the learner to find a model that strikes a balance between fitting the data and staying close to desired constraint weights. In this section, I describe the procedure for identifying the best compromise (i.e., the σ value at which the model is neither over- nor under-fitting the data) and present the results of this procedure.

In order to determine if the model is over- or under-fitting the data, authors such as Hayes and Wilson (2008) compare different grammars modeled on the training data to testing data from Wug tests (Berko, 1958) in which participants are asked to either produce or evaluate the acceptability of new data. As is customary in machine learning, the grammar that best predicts the test data after training on existing data is the one with the best fit. For studies that don't have access to new Wug data, however, cross-validation may serve as a substitute. While there are a number of cross-validation methods that vary slightly in their methodologies, all rely on a core premise: some part of the existing data set is designated for training and one or more other parts are designated as "new" test data.

Given the size of my dataset, I chose the Validation Set approach in which data is divided, at random, into two equal parts. In order to separate the dataset of 14,400 tableaux into two random subsets containing 7,200 tableaux each, I generated a random number to associate with each tableaux using the RAND() function in Excel, sorted the tableaux by those random numbers, and then designated the first half to be the training set and the second half to be the test, or validation, set. Subsequently, I fit various grammars to the training data by adjusting values for σ^2, beginning with 10,000 (the higher the σ value, the better the fit) and decreasing those values incrementally. I then tested each grammar on the validation set by entering in the constraint weights obtained in the various training models as default μs and using the resulting predictions to calculate the log likelihood[33] of each validation model. The σ that yielded the highest log likelihood for the validation data was identified as the value that achieved the best balance between data fit and faithfulness to default constraint weights. Table 17 compares the σ values used to train the original grammars with the log likelihoods of the validation models tested on those grammars.

Table 17

Validation data fit

σ^2 for all constraints	Log likelihood of validation (test) data
10000	-6783.21
1000	-6723.39
100	-6696.91
70	-6693.86
50	-6691.56
40	-6690.45
30	-6689.76
20	-6690.76
10	-6701.59
1	-6970.32

Here, we see that the log likelihood of the validation data is largest (i.e., closest to zero) when σ is set to 30 (although it should be noted that algorithm is not hypersensitive to the exact value of σ^2, so any value in the 20-50 range would likely yield similar results). Figure 45 shows a graphical representation of these values and compares them to the log likelihoods of the original training models.

[33] Log likelihood is a common measure of model fit. While log likelihood values are not useful in isolation as they are functions of sample size, they can be used comparatively to evaluate the fit of multiple models against one another. Log likelihood is calculated by multiplying the natural log of the predicted probability of each candidate by the number of actual observations and summing those products.

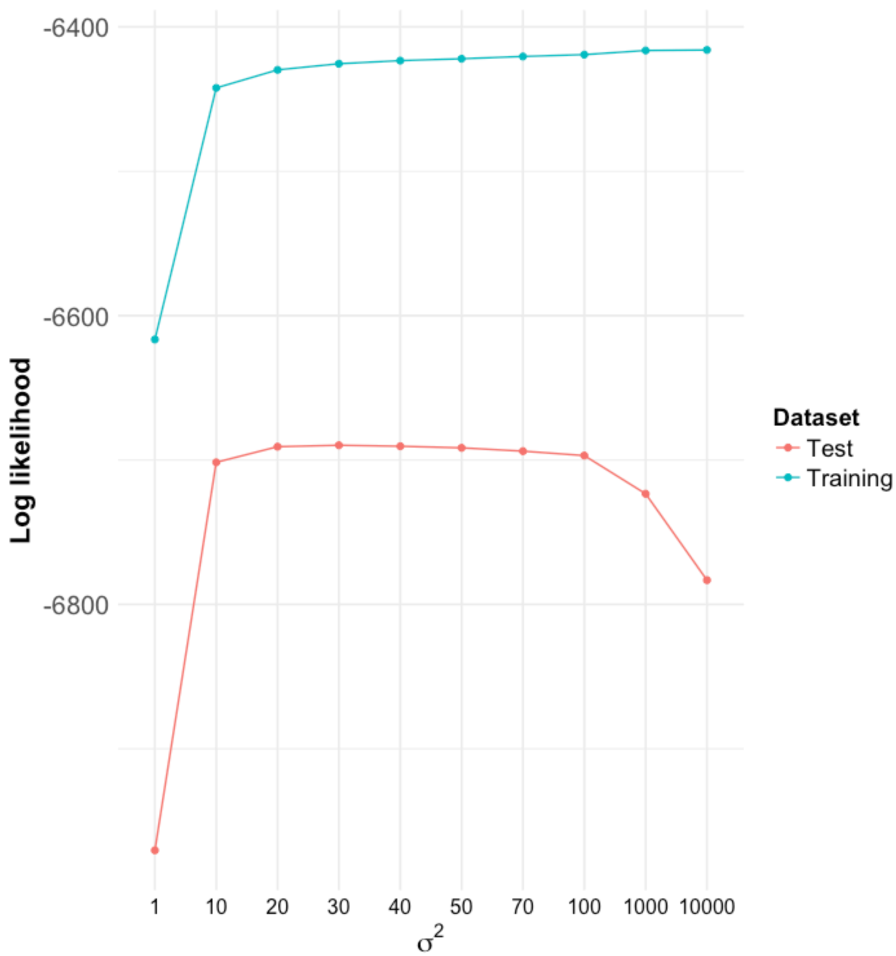

Figure 45. Comparison of log likelihoods for various models fitted to the training data and tested on the validation data.

In Figure 45 the reader can see that the log likelihood of the test data when σ^2 is large is relatively low, and then slowly improves as the value of σ decreases. At these higher σ^2 values, the model is overfitting because the algorithm has been told that fitting the observed data is more important than not deviating from designated constraint weights. As values of σ decrease, however, the balance shifts: the algorithm begins to care increasingly more about the constraint weights it has been told to strive for. This has a positive impact on fit until σ^2=30, at which point the log likelihood of the test data begins to decline because the algorithm has now strayed too far from the observed data. By the lowest value, σ^2=1, we see that the model is under-fitting both the training and test data.

The cross-validation procedure described in this section has identified σ^2=30 as the optimal parameter to mitigate both over- and under-fitting of the Salvadoran /s/ lenition data. The following section presents the base grammar for Salvadoran /s/ lenition and explores, in depth, what this grammar reveals about Markedness and Faithfulness within a phonetically-based, effort-based framework.

6.4 The base grammar

This section presents the base grammar for Salvadoran /s/ lenition; that is, a grammar that does not take the social characteristics of its speakers into account. Specifically, this section examines Markedness and Faithfulness constraints separately in order to evaluate my original hypotheses regarding the roles of effort cost and perceptual prominence in conditioning patterns of Salvadoran /s/ lenition.

6.4.1 Markedness in the grammar

Within Kirchner's (1998, 2004) effort-based framework for lenition, I have conceived of Markedness as an interaction between the inherent biomechanical effort required to produce a given gesture or group of gestures and the phonological context in which those sounds occur. According to this framework, the voiceless strident fricative [s] is particularly effortful because of the tight and sustained constriction it requires at the place of articulation as well as the glottal abduction needed to achieve voicelessness. I argue that, as a result of this effort cost, Salvadoran /s/ is frequently weakened to less effortful sounds by either reducing the degree of oral constriction, eliminating oral constriction altogether, forgoing glottal abduction, or any combination of these strategies. In order to explore this proposal, I hypothesized that more effortful sounds would be marked in phonological contexts that exacerbate their difficulty for phonetic reasons, giving way to easier yet less salient sounds, or even complete elision of /s/ when necessary. The following sections present Markedness constraint weights of the grammar by phonological context and discuss whether my original hypotheses have been confirmed or rejected.

The reader should note that all constraint weights discussed in this section represent a cumulative weight of some kind. That is, the constraint LAZY(any_gesture) was assigned a weight of 3.74 by the maxent learner, so weights for constraints LAZY(glottal_abduction), LAZY(oral_gesture), LAZY(strid_fric), and LAZY(vcl_strid_fric) represent the base penalty of 3.74 for producing any gesture plus additional penalties assigned for a particular constraint in a particular context. In some cases, the additional penalty is 0 and the resulting constraint weight is 3.74, indicating that the effort cost of that constraint in that context is the same as producing any gesture, regardless of type or context.

6.4.1.1 Intervocalic contexts

The first phonological context to be examined is intervocalic. Consistent with Kirchner's framework, I proposed that more open flanking segments would favor weakening of sounds that require tight oral constriction because the distance the active articulator must travel to reach the desired target increases. Therefore, I hypothesized that for sounds requiring some amount of oral constriction (i.e., [s], [z], and [$s^θ$]), lenition would be most prevalent between two vowels that are [-high] and least prevalent between two high vowels. For contexts in which one adjacent vowel is [-high] and the other [+high], I did not make any predictions apart from these contexts falling somewhere between $V_{[-high\ ]}$___$V_{[-high]}$ and $V_{[+high]}$ ___$V_{[+high]}$ with respect to effort cost. I also hypothesized that weights would be lowest in each of these contexts for Lazy(oral_gesture), but that there would be no difference for weights pertaining to Lazy(strid_fric) and Lazy(vcl_strid_fric). Figure 46 shows constraint weights learned for LAZY(vcd_strid_fric), LAZY(strid_fric), and LAZY(oral_gesture) in the four intervocalic contexts. Again, to be clear, all weights presented in the present section (6.4.1) are cumulative, as described on the previous page.

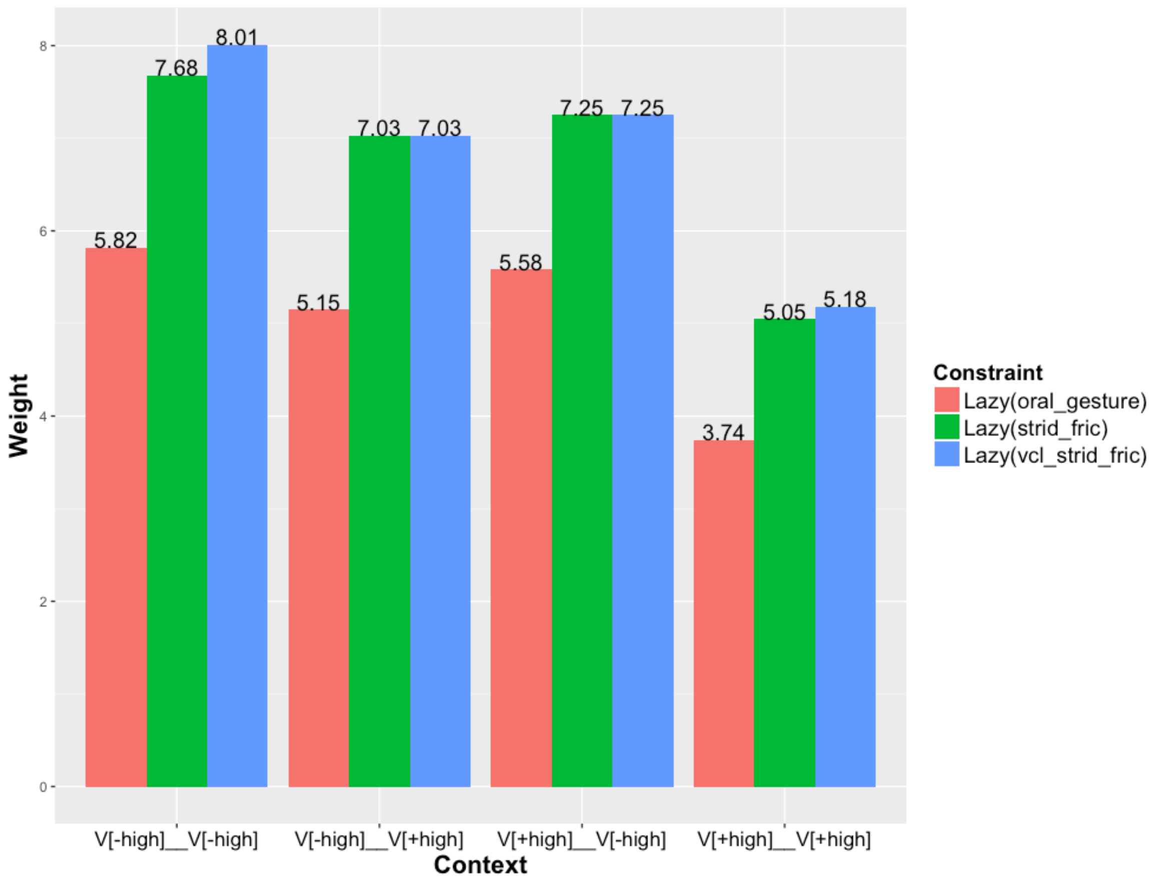

Figure 46. Constraint weights for LAZY(vcl_strid_fric), LAZY(strid_fric), and LAZY(oral_gesture) in the four intervocalic contexts.

Figure 46 reveals two important features of the grammar. First, regardless of the intervocalic context in which /s/ occurs, the impetus to avoid any strident fricatives ([s] or [z]) is greater than the impetus to avoid any oral gesture ([s], [z], or [s^0]). Second, we see that contextual hypotheses are confirmed: the need to "be lazy" is the greatest when /s/ occurs between two vowels that are [-high], and least when it occurs between two vowels that are [+high]. In fact, Lazy (oral_gesture) is weighted at 3.74 for $V_{[+high]}$ __ $V_{[+high]}$ contexts, suggesting that this context does not disfavor some lesser degree of oral constriction.

Furthermore, with respect to the $V_{[-high]}$ __ $V_{[+high]}$ and $V_{[+high]}$ __ $V_{[-high]}$ contexts, we see that all three constraints are ranked slightly higher in the latter context, suggesting that the openness of the following vowel might be slightly more important in determining the effort cost of a given context than the preceding vowel. Possible explanations for this difference will be discussed in Chapter 7.

While I hypothesized that vowel height would be an important impetus for lenition for sounds characterized by some degree of oral constriction, this factor should not have any impact on the constraint LAZY(glottal_abduction), whose goal is to reduce effort cost in the glottis, not the oral cavity. As predicted, weights for this constraint are more or less the same for all intervocalic contexts, as seen in Table 18 below.

Table 18

Constraint weights for LAZY(glottal_abduction) in intervocalic contexts

Constraint	Weight
Lazy(glottal abduction, V[-high]__V[-high])	5.40
Lazy(glottal abduction, V[-high]__V[+high])	5.49
Lazy(glottal abduction, V[+high]__V[-high])	5.41
Lazy(glottal abduction, V[+high]__V[+high])	5.68

6.4.1.2 Preconsonantal contexts

The second phonological environment to address are preconsonantal contexts, which were subcategorized according to the nature of the preceding segment (high vowel, non-high vowel, consonant) as well as the coronality and voicing of the following consonant. While predictions for these contexts did not follow from Kirchner, who only addresses intervocalic lenition, I hypothesized that the openness of the preceding segment as well as similarities in coronality and voicing between the class of allophone of /s/ and the following consonant would have implications for effort cost and therefore affect patterns of lenition. Specifically, I hypothesized that weights for LAZY(vcl_strid_fri) would be highest before consonants that are [-coronal] and [+voice] and lowest before consonants that are [+coronal] and [-voice], and then further conditioned by the height of the preceding vowel. I also hypothesized that constraint weights would be lowest for the C__C context due to the high degree of constriction of both the preceding and following segments. With respect to LAZY(strid_fric), I made the same predictions, but hypothesized that the impetus to "be lazy" would be lower preceding voiced consonants.

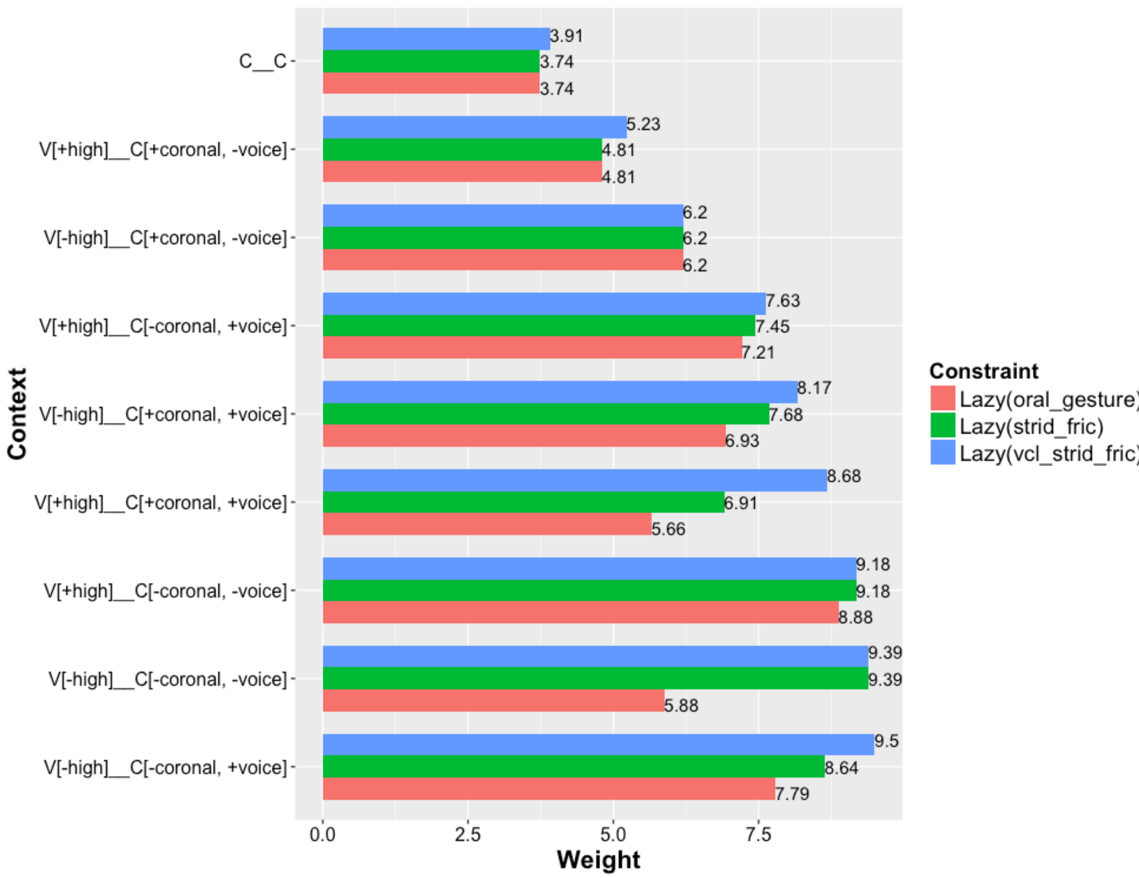

Figure 47. Constraint weights for LAZY(vcl_strid_fric), LAZY(strid_fric), and LAZY(oral_gesture) in preconsonantal contexts.

In Figure 47 we see that while the results for preconsonantal contexts are not as clear-cut as those for intervocalic contexts, many of my core hypotheses were confirmed. Consistent with my hypotheses are the findings that the effort cost of [s] is the highest in contexts in which the preceding vowel is more open and the following consonant does not agree with [s] in coronality or voicing. Furthermore, the C__C context is by far the easiest, with a weight approaching zero. The second easiest context for [s] is that in which the preceding vowel is high (i.e., more closed) and the preceding consonant agrees with [s] in coronality and voicing, and the third easiest context is the same context but with more open preceding vowels. While constraint weights for these four contexts are easily explained within an effort-based framework, the ordering of weights for the other five contexts is less clear. For example, it appears that coronality of the following segment plays the most important role, and that vowel height matters only when the following consonant is [-coronal]. Possible interpretations and explanations for these findings and will be discussed at length in Chapter 7.

Furthermore, as hypothesized, weights for LAZY(vcl_strid_fric) and LAZY(strid_fric) only differed when the following consonant was [+voice], confirming that the only factor differentiating the effort cost of [s] and [z] in preconsonantal contexts is the voicing of the following segment, which substantially decreases the effort cost of the latter. With respect to LAZY(oral_gesture), it appears that neither coronality nor preceding vowel height alone play a role in determining effort cost in preconsonantal contexts. In fact, the contexts in which weights for this constraint differ most dramatically from those of the other two are $V_{[-high]}$__$C_{[-coronal, -voice]}$ and $V_{[+high]}$__$C_{[+coronal, +voice]}$, suggesting that perhaps there is something particularly easy about executing an oral gesture in these specific contexts. These findings will be further explored in the next chapter.

With respect to LAZY(glottal_abduction) in preconsonantal contexts, I hypothesized that voicing of the following consonant would have a significant effect on the effort cost of glottal abduction. That is, when preceding a consonant that is [+voice], the effort cost to maintain voicelessness via an open glottis increases and therefore the impetus for /s/ to "be lazy" in those conditions should be higher than when it precedes another voiceless consonant. Figure 48 below shows the weights found by the maxent learner for LAZY(glottal_abduction) in these contexts.

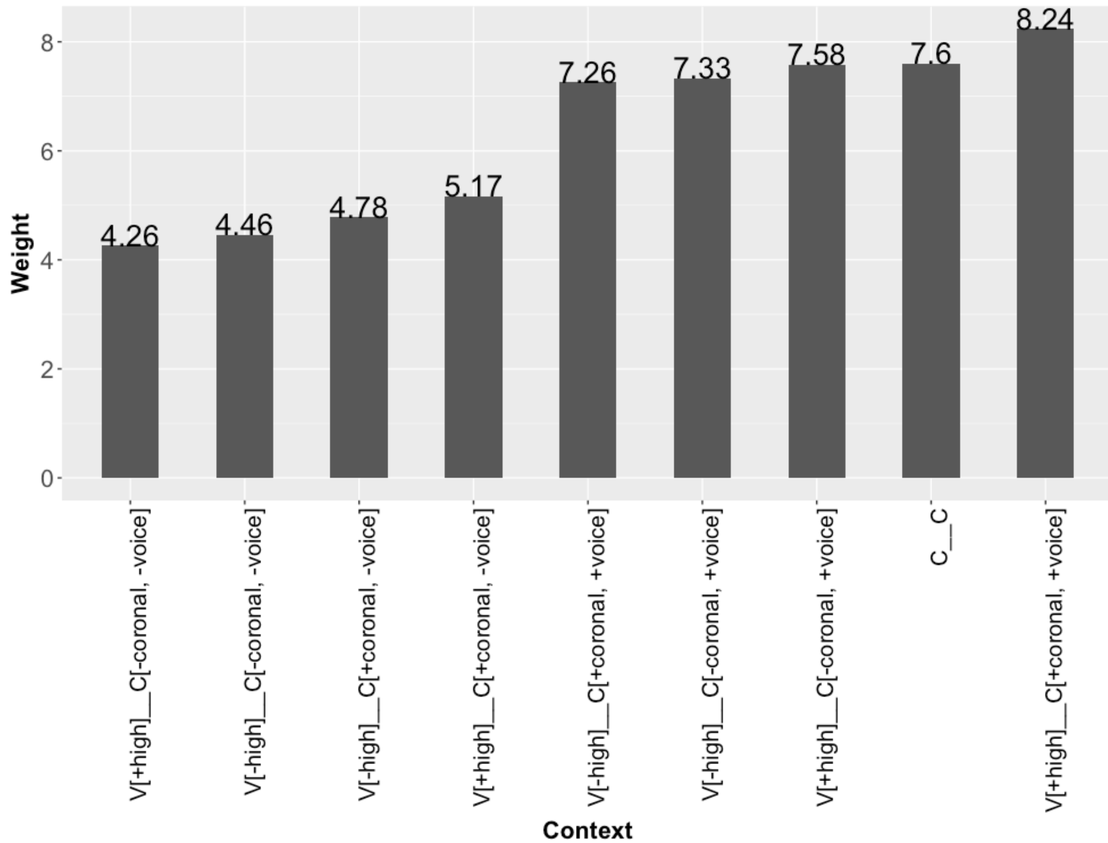

Figure 48. Constraint weights for LAZY(glottal_abduction) in preconsonantal contexts.

As the reader can see, the hypothesis with respect to LAZY(glottal_abduction) in preconsonantal contexts was strongly confirmed: constraint weights are notably higher when the following consonant is [+voice]. The only curious finding is the high ranking of C__C in which /s/ always occurs before a voiceless consonant (/t/ or /p/). The four possible C__C environments for /s/ in the dataset are /ksp/, as in the word *explicar* /eksplikaɾ/ 'to explain,' /kst/, as in the word *sexto* /seksto/ 'sixth,' /nst/, as in the word *construir,* /konstruiɾ/ 'to build,' and /nsp/, as in *transparente* /transpaɾente/. Therefore, I drilled down into these contexts (n = 27) to see if the voicing of the preceding consonant (i.e., /n/ vs. /k/) might be driving this high constraint weight.

Figure 49 compares the proportion of allophones requiring glottal abduction in /ks/ clusters with the same proportions in /ns/ clusters.

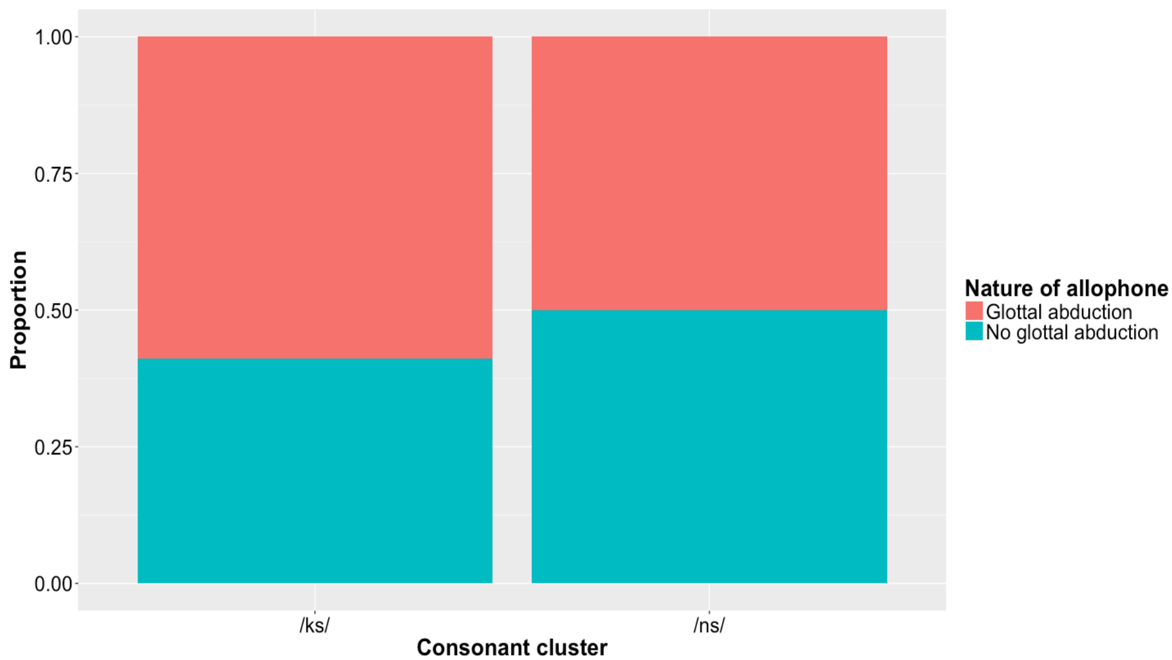

Figure 49. Comparison of the proportion of allophones requiring glottal abduction in /ks/ clusters (n = 17) with the proportion of allophones requiring glottal abduction in /ns/ clusters (n = 10).

While Figure 49 shows that allophones requiring glottal abduction are slightly more common in the /ks/ context (58.82%) than in the /ns/ context (50%), these differences are likely not significant given the sample size. Alternative explanations for the relatively high ranking of LAZY(glottal_abduction) in C__C contexts will be explored in Chapter 7.

6.4.1.3 C__V contexts

The third phonological environment to address is the C__V context, which was again subcategorized according to the height of the following vowel. I hypothesized that, for LAZY (vcl_strid_fric), LAZY(strid_fric), LAZY(oral_gesture), a lower following vowel would favor lenition, and that for LAZY(glottal_abduction) the height of the following vowel wouldn't have any effect. Figure 50 shows the weights for all four constraints in this context.

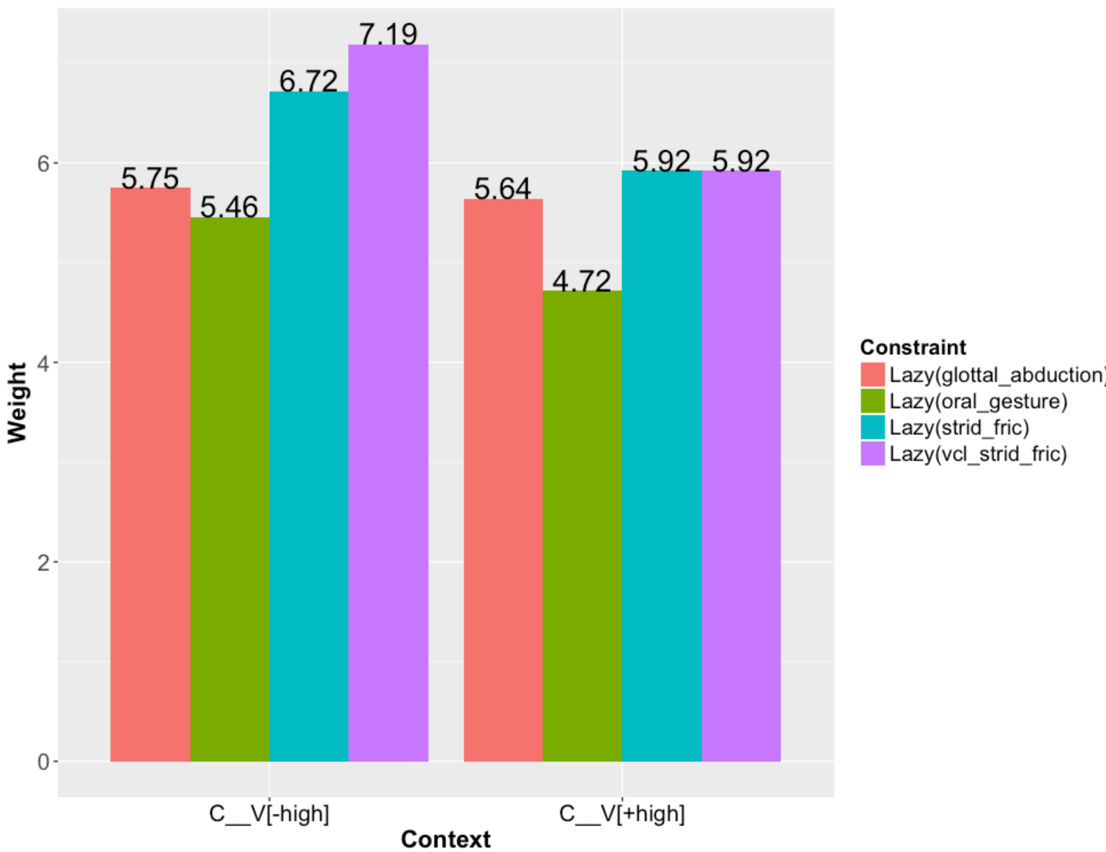

Figure 50. Constraint weights for LAZY (vcl_strid_fric), LAZY(strid_fric), LAZY(oral_gesture), and LAZY(glottal_abduction) in C__V contexts.

Consistent with my hypotheses, vowels that are [-high] favor lenition with respect to oral constriction, and glottal abduction is essentially unaffected by following vowel height.

Furthermore, while I hypothesized that constraint weights for LAZY (vcl_strid_fric) and LAZY(strid_fric) would not differ in C__V contexts, it appears that the former is weighted more heavily in the C__$V_{[-high]}$ context. It is possible that this is due to differences in voicing of the preceding consonant, as we saw in Figure 49.

6.4.1.4 Pause-adjacent contexts

The fourth and final phonological environment to examine is pause-adjacent, which was subdivided based on the location of the pause (before /s/ or after /s/) as well as the height of the adjacent vowel. I hypothesized that pause-adjacent contexts would be weighted according to the height of the adjacent (preceding or following) vowel for the three constraints that penalize various degrees of oral constriction. I did not propose any hypotheses with respect to LAZY(glottal_abduction). Figure 51 shows the weights found by the maxent learner for all four constraints in pause-adjacent contexts.

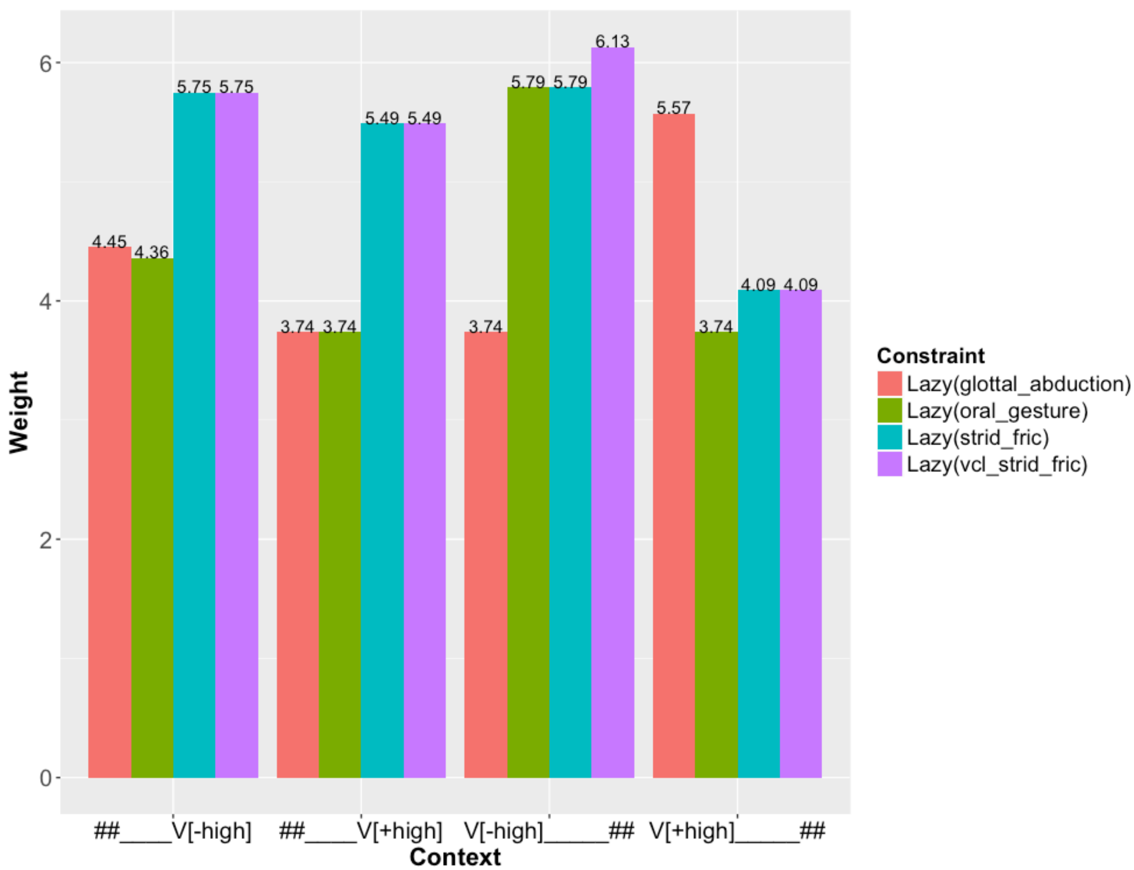

Figure 51. Constraint weights for LAZY(vcl_strid_fric), LAZY(strid_fric), and LAZY(oral_gesture) LAZY(glottal_abduction) in pause-adjacent contexts.

As the reader can see, the height of the adjacent vowel has a systematic effect on constraint weights for LAZY(vcd_strid_fric), LAZY(strid_fric), and LAZY(oral_gesture) in pause-adjacent contexts. Weights for LAZY(glottal_abduction), contrastingly, depend on the location of the adjacent pause, increasing on average in pre-pausal contexts. As always, vowel height is not meaningful for this constraint type.

6.4.1.5 Markedness: Summary

This section has presented findings related to Markedness in the maxent grammar. Because this dissertation proposes that effort-based lenition can, in part, be accounted for via the interaction between gesture types and phonological environments, I have presented results context by context. The highest-ranking Markedness constraints in the grammar incentivize various reductions to the degree of oral constriction preceding consonants that are [-coronal], adjacent to vowels that are [-high], or both. Lenition in the form of reduced glottal abduction, on the other hand, is necessitated before consonants that are [+voice] and in C__C contexts. The lowest-ranking Markedness constraints in the grammar indicate that oral constriction is least costly before consonants that match /s/ in both coronality and voicing, adjacent to vowels that are [+high], and in C__C contexts; glottal abduction is most permissible after a pause or preceding a voiceless consonant. An interpretation and discussion of these results—both those that confirm the original hypotheses and those that no not—is offered in Chapter 7.

6.4.2 Faithfulness in the grammar

Following Kirchner's (1998, 2004) framework, I have conceived of Faithfulness as the need to preserve particular features of /s/ in particular prosodic positions in order to maintain important perceptual distinctions. These constraints interact with the Markedness constraints discussed in the previous sections to render the observed patterns of Salvadoran /s/ lenition. I styled my Faithfulness constraints in the form 'PRESERVE(feature)/ prosodic position, tonicity,' hypothesizing that stronger prosodic positions would favor stronger segments, and that those differences would be further modulated by the tonicity of the syllable containing /s/. The following sections present learned Faithfulness constraint weights by feature and discuss whether the original hypotheses have been confirmed or rejected.

6.4.2.1 PRESERVE(strident, voice)

The most fundamental Faithfulness constraints in the grammar are those pertaining to PRESERVE(strident, voice), which requires that both input [+strident] and [-voice] specifications be maintained in the output, therefore penalizing all non-[s] allophones. I hypothesized that the grammar would prioritize the preservation of these features according to prosodic strength relations, and that within each prosodic position, weights would be slightly higher in tonic syllables than in atonic ones. Figure 52 shows the weights learned by the maxent grammar for this set of constraints.

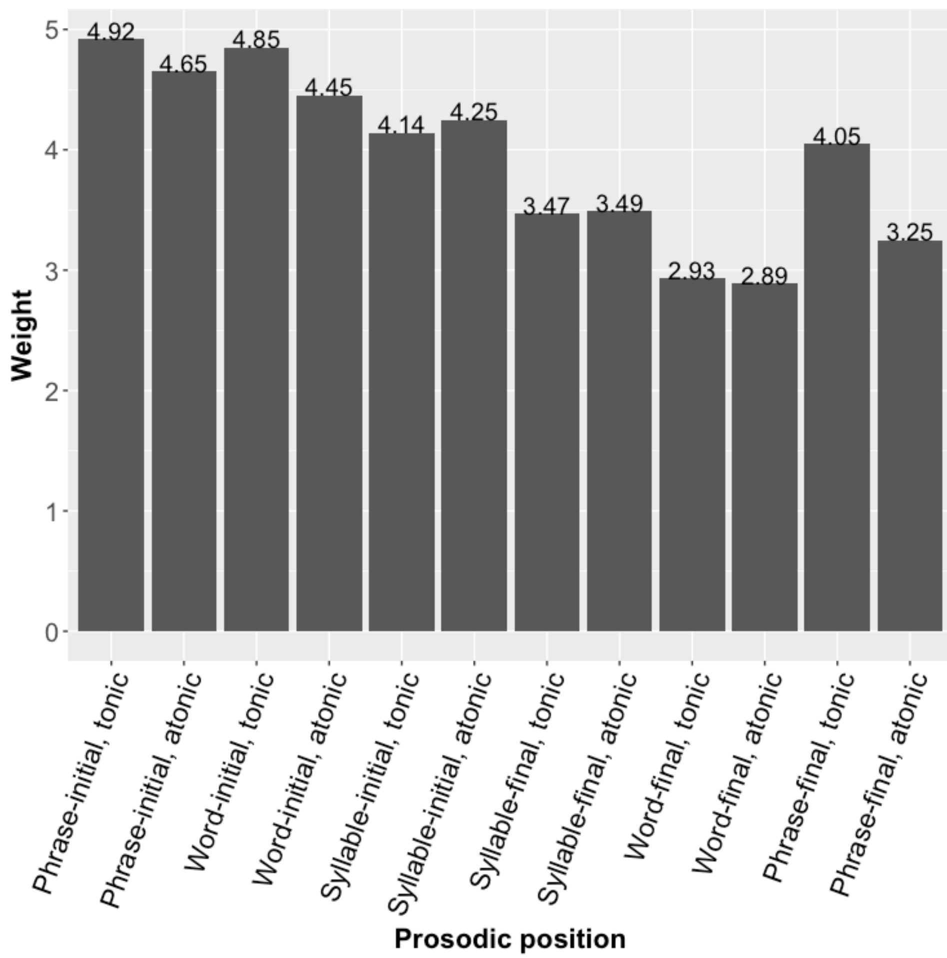

Figure 52. Constraint weights for PRES(strid, voice) in six prosodic positions, further divided according to syllable tonicity.

As the reader can see, the most fundamental hypothesis was confirmed: the need to preserve input specifications for the features [+strident] and [-voice] is greatest in the strongest prosodic positions and least in the weakest. While predictions regarding phrase-final position were difficult to hypothesize, we see an effect of phrase-final lengthening in that PRESERVE(strident, voice) is weighted higher in phrase-final position than in word-final, phrase-medial position. As predicted, syllable tonicity plays a role in determining relative constraint weights; however, this difference does not appear to be important word-medially (compare syllable-initial tonic vs. atonic and syllable-final tonic vs. atonic).

The sections that follow compare the weights learned for PRESERVE(strident, voice) with those learned for the four more general PRESERVE constraints. These sections indicate which features are most and least important to preserve in the output for the twelve prosodic positions in question, thereby revealing the forces driving the more nuanced differences in lenition patterns beyond [s] (most faithful yet most marked) vs. [∅] (least marked yet least faithful).

6.4.2.2 PRESERVE(strident)

The first subgroup of PRESERVE constraints to examine are those pertaining to PRES(strident), which penalizes outputs in which the input [+strident] specification has not been maintained and is a key differentiator between allophones [z] and [s^{θ}].

Figure 53 below compares weights for PRES(strid, voice) and PRES(strid) in order to highlight where the latter assigns additional penalties and where it patterns similarly to the former.

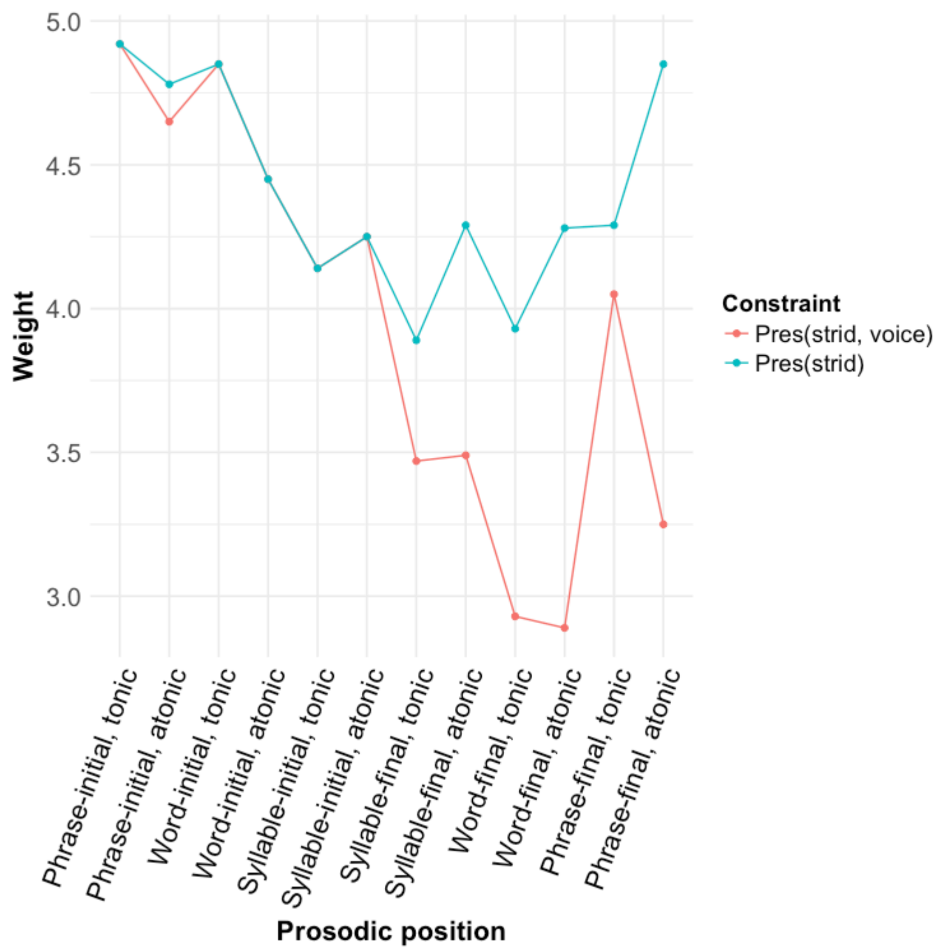

Figure 53. Comparison of constraint weights for PRESERVE(strident, voice) and PRESERVE(strident) in 12 prosodic positions.

In Figure 53, the reader can see that PRESERVE(strident) assigns almost no additional weights in stronger prosodic positions, suggesting that [z] and [s^{θ}] are equally "bad" with respect to Faithfulness. The prosodic positions in which PRESERVE(strident) assigns large additional penalties—syllable-final, word-final, and phrase-final positions—are the contexts in which [s^{θ}] is worse than [z]. These issues will be explored in more depth in the next chapter.

6.4.2.3 PRESERVE(voice)

The next subgroup of PRESERVE constraints to examine are those pertaining to PRESERVE(voice), which penalizes outputs in which the input [-voice] specification has not been maintained and is a key differentiator between allophones [s^{θ}] and [z] as well as [h] and [ɦ]. Figure 54 below compares weights for PRESERVE(strident, voice) and PRESERVE(voice) in order to highlight where the latter assigns additional penalties and where it patterns similarly to the former.

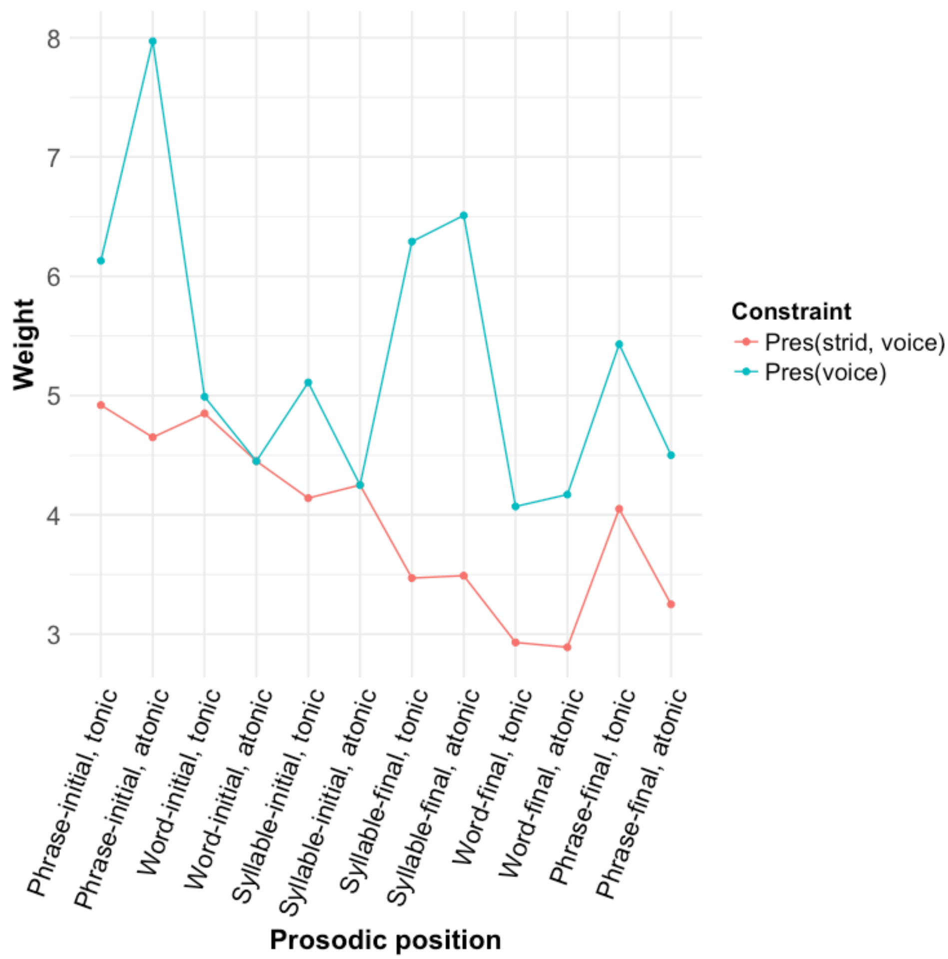

Figure 54. Comparison of constraint weights for PRES(strid, voice) and PRES(voice) in 12 prosodic positions.

Figure 54 shows that the constraint PRESERVE(voice) patterns most differently from PRESERVE(strident, voice) in phrase-initial and syllable-final positions. Because debuccalization rarely occurs phrase-initially, this finding likely suggests that [z] is dispreferred in this strongest prosodic position as compared to [s^{θ}]. The large additional penalty for not preserving the input [-voice] specification syllable-finally likely suggests that [h] is strongly preferred to [ɦ] in this position. This is logical, as syllable-final is the strongest of the weaker set of prosodic positions and therefore prefers the stronger of the two debuccalized allophones, the voiceless glottal fricative.

6.4.2.4 PRESERVE(coronal)

The next subgroup of PRESERVE constraints to examine are those pertaining to PRESERVE(coronal), which penalizes outputs in which the input [+coronal] specification has not been maintained and is a key differentiator between allophones with and without some degree of oral constriction. Figure 55 compares weights for PRESERVE(strident, voice) and PRESERVE(coronal) in order to highlight where the latter assigns the highest additional penalties.

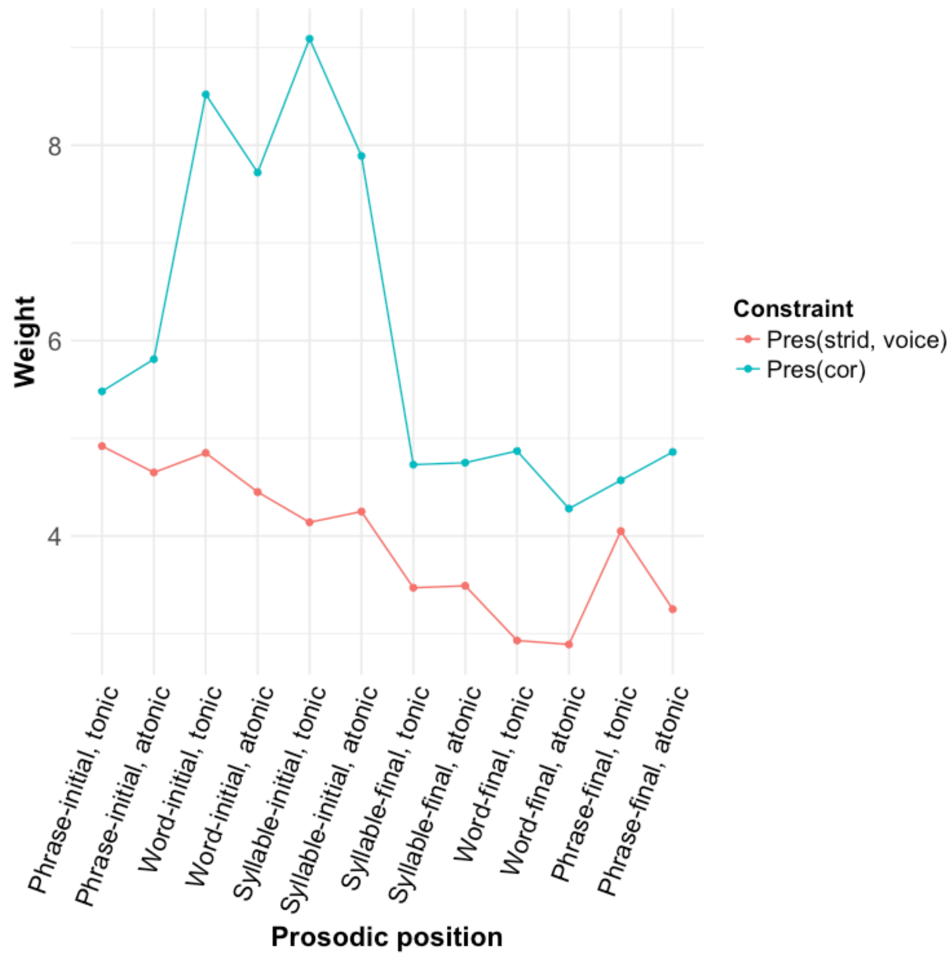

Figure 55. Comparison of constraint weights for PRESERVE(strident, voice) and PRESERVE(coronal) in 12 prosodic positions.

In Figure 55, the red line indicates penalties for using any lenited allophone—be it [z], [s^{θ}], [h], [ɦ], or [Ø]—while the blue line indicates additional penalties for using one of the three non-coronal allophones, [h], [ɦ], or [Ø]. In other words, larger differences between the two constraints highlight prosodic positions in which having a consonant with no place of articulation is much worse than having a consonant that is simply not [s]. This difference is not very large in phrase-initial position because PRESERVE(strident, voice) is already highly ranked, indicating that the more important distinction in this position is [s] vs. other allophone, not oral constriction vs. no oral constriction.

Differences are also smaller in the three coda positions, but for a different reason: maintaining the [+coronal] specification here is less important because these prosodic positions are less salient. Loss of the [+coronal] input specification is most heavily penalized word- and syllable-initially, positions in which [z] and [s^{θ}] are acceptable substitutes for [s] because of their perceptual salience, but weaker [h], [ɦ], and [∅] are not. Furthermore, in contrast to the additional penalties incurred for PRESERVE(strident) and PRESERVE(voice), there is some increase in weight for each of the 12 PRESERVE(coronal) constraints.

6.4.2.5 PRESERVE(segmental)

The final subgroup of PRESERVE constraints to examine are those pertaining to PRESERVE(segmental), which penalizes outputs in which the input [+segmental] specification has not been maintained and is a key differentiator between [ɦ] and [∅]. Figure 56 compares weights for PRESERVE(strident, voice) and PRESERVE(segmental) in order to highlight where the latter assigns the highest additional penalties.

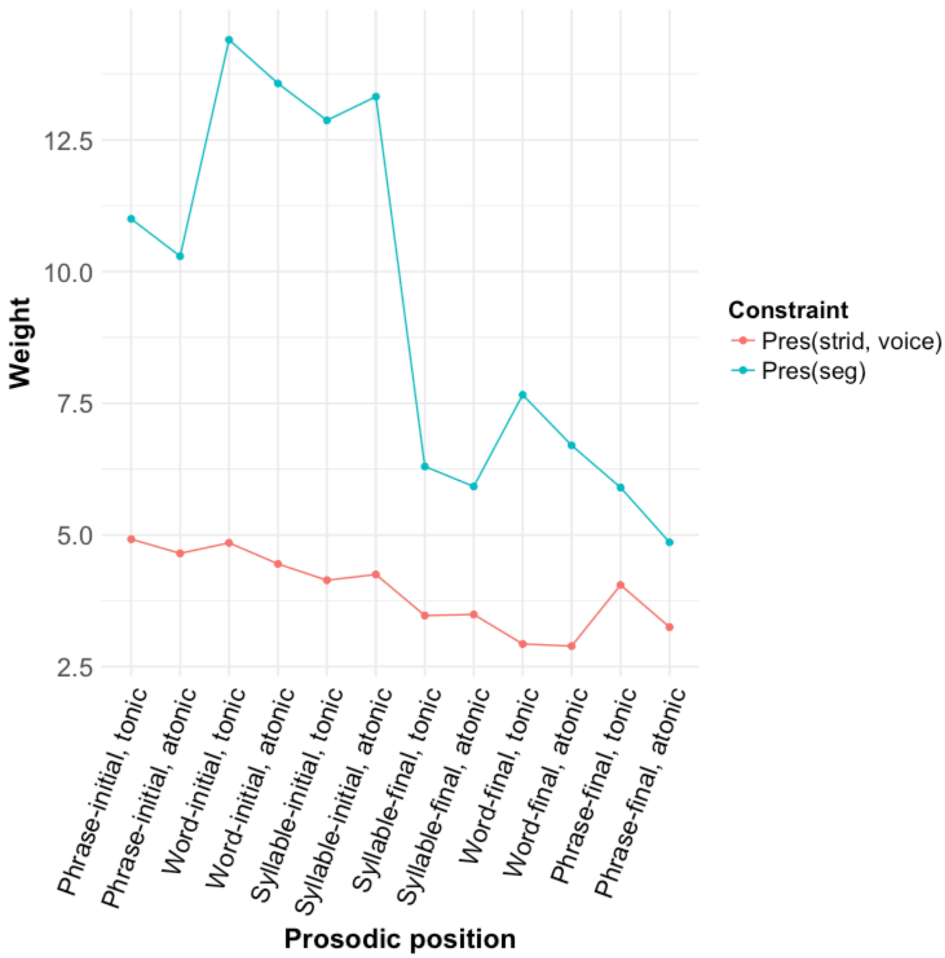

Figure 56. Comparison of constraint weights for PRESERVE(strident, voice) and PRESERVE(segmental) in 12 prosodic positions.

Notably, the PRESERVE(segmental) contour in Figure 56 appears to parallel the original PRESERVE(strident, voice) contour the most closely of any constraint group, albeit at much more extreme weights for every prosodic position. As with the original contour, syllable tonicity plays a role in distinguishing weights within prosodic categories except when /s/ occurs word-medially. Additionally, similar to the PRESERVE(coronal) contour, we see that the most dramatic penalties for PRESERVE(segmental) occur word-initially and word-medially, while those in phrase-initial position remain slightly lower. This is, again, likely because there is so little lenition phrase-initially as a whole that the difference between weights for PRESERVE(strident, voice) and PRESERVE(segmental) can be slightly smaller.

6.4.2.6 Faithfulness: Summary

This section has presented findings related to Faithfulness in the maxent grammar. Because this dissertation proposes that /s/ lenition-blocking can, in part, be accounted for via constraints that discourage the loss of particular features in particular prosodic positions, I have presented results feature by feature. The highest-ranking Faithfulness constraints in the grammar require the preservation of input [+segmental] and [+coronal] features—those that provide the most fundamental distinctions between [s] and non-[s] allophones—in more salient prosodic positions. Constraint weights affecting input specifications [-voice] and [+strident]—which primarily serve to differentiate between allophones that have some degree of oral constriction—present a more nuanced picture of Salvadoran /s/ weakening in which these feature changes are blocked in particular subsets of prosodic positions. Finally, consistent with original hypotheses, constraint weights necessitating the preservation of both [-voice] and [+strident] specifications—which taken together differentiate [s] from all other possible outputs—are highest in more salient prosodic positions and are further modulated according to syllable tonicity. An interpretation and discussion of these results, beyond what has been offered here, can be found in Chapter 7.

Sections 6.4.1 and 6.4.2 have served to separately explore the roles of context-based Markedness and position-based Faithfulness in the maxent grammar. Now, in the following section, we turn to the interaction between Markedness and Faithfulness in conditioning patterns of Salvadoran /s/ lenition.

6.4.3 The interaction between Markedness and Faithfulness

This section will explore the interaction between Markedness and Faithfulness in the /s/ lenition grammar via a series of illustrative tableaux. The primary goal of this section is to examine the role of various constraints in determining predicted output probabilities and thereby identify the most powerful forces at work in the grammar. The following subsections are divided in such a way that important distinctions between allophones are as clear as possible.

The reader should note that constraint weights displayed in this section represent *additional* penalties, not summed penalties as were presented in the previous section. For example, while the total penalty for LAZY(vcl_strid_fric, $V_{[-high]}$__$V_{[-high]}$) is 8.01, this number is broken down into its various parts in the forthcoming tableaux: a penalty of 3.74 for LAZY(any_gesture), an additional penalty of 2.08 for LAZY(oral_gesture), an additional penalty of 1.86 for LAZY(strident_fricative), and an additional penalty of 0.34 for LAZY(vcl_strid_fric). The choice to present weights in this way is for illustrative purposes, as both the calculation of candidates' harmonies as well as the power of individual constraints in the grammar will be more transparent.

6.4.3.1 Configurations favoring oral constriction

Let's first explore contexts that favor varying degrees of oral constriction, beginning with inputs for which [s] is almost the sure output candidate. We begin by examining the word *siempre* /sjempre/ ('always'). The /s/ in *siempre* always corresponds to one of two strong prosodic positions, word-initial or phrase-initial, and carries lexical stress. Given these conditions, Faithfulness constraints at work for this input will always be highly ranked, favoring retention of [s] and disfavoring weaker outputs. How probable [s] is compared to other potential outputs, however, depends on the effort cost of the context in which /s/ occurs. The more difficult the context, the more likely it is that Markedness will trump Faithfulness, rendering a non-[s] allophone in the output. Tables 19A-D show four tableaux for /sjempɾe/ by varying contexts. The first column displays the input and possible output forms, columns two through 11 show constraints, weights for the pertinent phonological context or prosodic position, and violations incurred by candidates. The last three columns show a candidate's harmony (a sum of constraint violations for a given constraint times its weight), the probability the grammar assigns to this candidate, and the number of actual observations of this candidate in the data.

Tables 19a-19d

Tableaux for /sjempɾe/ (‘always’)

a. Input /sjempɾe/ where /s/ follows a pause

##/sjempɾe/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.74	0.00	0	1.75	0.00	4.92	0.00	0.56	1.21	5.51	Harmony	Predicted	Observed
a. ˈsjem.pɾe	*	*	*	*	*						5.49	0.93	27
b. ˈzjem.pɾe	*		*	*		*			*		11.62	0.00	0
c. ˈs$^{\theta}$jem.pɾe	*	*	*			*	*				8.66	0.04	4
d. ˈhjem.pɾe	*	*				*	*	*			9.22	0.02	0
e. ˈɦjem.pɾe	*					*	*	*	*		10.43	0.01	0
f. ˈØjem.pɾe						*	*	*	*	*	12.20	0.00	0

b. Input /sjempɾe/ where /s/ follows a consonant

/d, l, m, n, ɾ # sjempɾe/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.74	1.90	0.98	1.20	0.00	4.85	0.00	3.67	0.14	5.88	Harmony	Predicted	Observed
a. ˈsjem.pɾe	*	*	*	*	*						7.82	0.92	8
b. ˈzjem.pɾe	*		*	*		*			*		10.91	0.04	0
c. ˈs$^{\theta}$jem.pɾe	*	*	*			*	*				11.47	0.02	0
d. ˈhjem.pɾe	*	*				*	*	*			14.16	0.00	0
e. ˈɦjem.pɾe	*					*	*	*	*		12.40	0.01	0
f. ˈØjem.pɾe						*	*	*	*	*	14.54	0.00	0

c. Input /sjempɾe/ where /s/ follows a high vowel

/i, u # sjempɾe/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.74	1.94	0.00	1.31	0.13	4.85	0.00	3.67	0.14	5.88			
a. ˈsjem.pɾe	*	*	*	*	*						7.12	0.91	3
b. ˈzjem.pɾe	*		*	*		*			*		10.04	0.05	1
c. ˈs$^{\theta}$jem.pɾe	*	*	*			*	*				10.53	0.03	1
d. ˈhjem.pɾe	*	*				*	*	*			14.20	0.00	0
e. ˈɦjem.pɾe	*					*	*	*	*		12.40	0.00	0
f. ˈØjem.pɾe						*	*	*	*	*	14.54	0.00	0

d. Input /sjempɾe/ where /s/ follows a vowel that is [-high]

/e, o, a # sjempɾe/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.74	1.75	1.41	1.88	0.00	4.85	0.00	3.67	0.14	5.88			
a. ˈsjem.pɾe	*	*	*	*	*						8.78	0.89	19
b. ˈzjem.pɾe	*		*	*		*			*		12.02	0.03	0
c. ˈs$^{\theta}$jem.pɾe	*	*	*			*	*				11.75	0.05	0
d. ˈhjem.pɾe	*	*				*	*	*			14.01	0.00	0
e. ˈɦjem.pɾe	*					*	*	*	*		12.40	0.02	0
f. ˈØjem.pɾe						*	*	*	*	*	14.54	0.00	0

In Tables 19a-d, [s] is the predicted winner at a chance of 89% or more because, as the reader can see, Faithfulness in these prosodic positions is weighted heavily. Even though the observed outputs are the most diverse in Tables 19a, with 27 observations of [s] and 4 observations of [s^{θ}], candidate ##[s]iempre still receives the greatest share of the harmony in this tableau as compared to Tables b-d because the impetus to reduce effort cost between a pause and a high vowel is the lowest. Then, as the degree of constriction slowly increases from a preceding consonant in Tables 19d to a preceding [-high] vowel in Tables 19d, the probability of output [s] incrementally decreases, and a greater percentage of the probability is shared with other candidates.

Tables 19a-d also reveal some important observations with respect to non-[s] allophones. In Tables 19a, outputs [s^{θ}] and [h] are the second and third most probable candidates because preserving voicelessness in phrase-initial position is so important that it trumps the preservation of coronality. Output [ˈs^{θ}jem.pɾe] is twice as probable than output [ˈhjem.pɾe], however, due to the penalty for not preserving the input [+coronal] specification. In this tableau, even [ɦ] is slightly more harmonic than [z]: while neither of them do well on Faithfulness because of the perceptual cost of voicing in this prosodic position, the former is easier and therefore considered a slightly better candidate by the grammar.

In Tables 19b-d, however, [z] and [s^{θ}] take the lion's share of the non-[s] probability because of the importance of preserving the input [+coronal] specification in word-initial position. In Tables 19b-c, [z] is more harmonic than [s^{θ}] because of the comparatively lower effort cost of producing a strident in these more constricted contexts. However, by Tables 19d, the effort cost associated with producing a strident fricative in the more open $V_{[-high]}$ __$V_{[+high]}$ context renders the gesturally-undershot [s^{θ}] the more harmonic candidate.

Let's now turn our attention entirely to outputs [z] and [s^θ], which present a particularly interesting case in the grammar. These allophones—neither of which is part of the /s/ lenition paradigm in many dialects of Spanish—are never chosen as the most harmonic candidate in any tableau. Within an effort-based framework, I argue that this is because these intermediate allophones are more effortful than those that require no oral constriction but are worse than [s] with respect to Faithfulness. In fact, the number of occurrences of each allophone in the raw data reflects exactly this: the "best" candidates in the grammar are those at the extreme ends of the lenition spectrum, i.e., [s] (the most faithful) and [∅] (the least effortful). Outputs [z] and [s^θ], then, surface when the speaker is striving for [s] in stronger prosodic positions but misses the target due to some context-conditioned increase in effort cost.

While candidates [z] and [s^θ] are never the most harmonic, the maxent grammar still assigns some probability to each of the output candidates, allowing us to examine the contexts in which these allophones are the most probable. Table 20 shows the tableau for the word *pienso* /pjenso/ ('I think'), which represents the most probable context for [z]: /s/ occurs in a C__V context, which is not difficult with respect to oral constriction but presents a challenge with respect to glottal abduction as the preceding consonant is voiced 95% of the time.[34] In this example, /s/ occurs in an atonic, quasi-strong prosodic positon (syllable-initial). This tableau is juxtaposed with that presented in Table 21, which shows the tableau for the word *salir* /saliɾ/ ('to leave'), whose /s/ also occurs in C_V context but this time in word-initial, atonic prosodic position, representing the second-most amenable context for [z] in the base grammar. The preceding consonant in this context was voiced 100% of the time yet the probability of [z] decreases slightly as the strength of the prosodic position increases, favoring [s].

[34] In word-medial, C__V contexts, the preceding consonant was [+voice] 929 times and [-voice] 47 times.

Table 20

Tableau for /pjenso/ ('think'; 1st person singular)

/pjenso/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.74	2.01	1.72	1.26	0.47	4.25	0.00	3.64	0.00	5.44	Harmony	Predicted	Observed
a. ˈpjen.so	*	*	*	*	*						9.20	0.73	15
b. ˈpjen.zo	*		*	*		*			*		10.97	0.12	3
c. ˈpjen.s^{θ}o	*	*	*			*	*				11.72	0.06	0
d. ˈpjen.ho	*	*				*	*	*			13.64	0.01	0
e. ˈpjen.ɦo	*					*	*	*	*		11.63	0.06	1
f. ˈpje.nØo						*	*	*	*	*	13.33	0.01	0

Table 21

Tableau for /saliɾ/ ('to leave')

/b, k, l, n, p, ɾ, t #saliɾ/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.74	2.01	1.72	1.26	0.47	4.45	0.00	3.27	0.00	5.85	Harmony	Predicted	Observed
a. saˈliɾ	*	*	*	*	*						9.20	0.75	3
b. zaˈliɾ	*		*	*		*			*		11.17	0.10	0
c. s^{θ}aˈliɾ	*	*	*			*	*				11.92	0.05	0
d. haˈliɾ	*	*				*	*	*			13.47	0.01	0
e. ɦaˈliɾ	*					*	*	*	*		11.46	0.08	0
f. Øaˈliɾ						*	*	*	*	*	13.57	0.01	0

Table 20 represents the context in which [z] is most likely to occur. While [s] is still the most harmonic output because it doesn't violate any Faithfulness constraints, [z] is the second-most harmonic in a context in which glottal abduction is particularly challenging. In this example, candidates [ˈpjen.s$^{\theta}$o] and [ˈpjen.ɦo] are equally probable: the former does much worse with respect to effort cost, yet the former is heavily penalized with respect to Faithfulness. However, in Table 21, we see probabilities shift slightly. In the example /b, k, l, n, p, ɾ, t #saliɾ/, ouput [saˈliɾ] is slightly more probable and output [zaˈliɾ] is slightly less probable than in Table 20 due to the higher ranking of PRESERVE(strident, voice) in word-initial position. However, output [ɦaˈliɾ] is now slightly more probable than [s$^{\theta}$aˈliɾ] because LAZY(glottal_abduction) is relatively more important than PRESERVE(cororonal) as compared to Table 20.

The examples above demonstrate that [z] is favored in atonic, word- and syllable-initial positions in C__V contexts. Contrastingly, as can be seen in Table 22—which displays the tableau for the word *pasar* /pasaɾ/ ('to pass') — [s$^{\theta}$] is preferred primarily in tonic, syllable-onset position in intervocalic contexts. Table 23 displays the tableau for the word *casa* /kasa/ ('house'), which represents the second-most amenable context for this allophone: intervocalic in word-medial, atonic position.

Table 22

Tableau for /pasaɾ/ (‘to pass’)

/pasaɾ/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesure)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.74	1.66	2.08	1.86	0.37	4.14	0.00	4.95	0.97	3.78			
a. paˈsaɾ	*	*	*	*	*						9.71	0.82	23
b. paˈzaɾ	*		*	*		*			*		12.79	0.04	0
c. paˈs$^{\theta}$ɾ	*	*	*			*	*				11.62	0.12	3
d. paˈhaɾ	*	*				*	*	*			14.49	0.01	0
e. paˈɦaɾ	*					*	*	*	*		13.80	0.01	0
f. paˈØaɾ						*	*	*	*	*	13.84	0.01	0

Table 23

Tableau for /kasa/ (‘house’)

/kasa/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesure)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.74	1.66	2.08	1.86	0.37	4.25	0.00	3.64	0	5.44			
a. ˈka.sa	*	*	*	*	*						9.71	0.70	30
b. ˈka.za	*		*	*		*			*		11.93	0.07	0
c. ˈka.s$^{\theta}$a	*	*	*			*	*				11.73	0.09	11
d. ˈka.ha	*	*				*	*	*			13.29	0.02	0
e. ˈka.ɦa	*					*	*	*	*		11.63	0.10	4
f. ˈkaØa						*	*	*	*	*	13.33	0.02	0

When we compare the tableaux above, we see that both [s] and [s^{θ}] are more probable in the former than the latter due to the higher values for both PRESERVE(coronal)—a crucial differentiator between [s^{θ}] and [h], [ɦ]—and PRESERVE(voice), a crucial differentiator between [s^{θ}] and [z], [ɦ]. Because these constraints are more highly weighted in the tonic position, [s] and [s^{θ}] are by far the most viable candidates in Table 22. However, in the same context but in atonic position as shown in Table 23, weights for both PRESERVE(coronal) and PRESERVE(voice) are reduced and, as a result, [ˈka.s^{θ}a] is slightly less probable than [paˈs^{θ}ɾ]. In fact, because of the demotion of Faithfulness in atonic, syllable-initial position, [ˈka.ɦa] is now the second-most likely candidate: it has low effort cost, and does significantly better on Faithfulness than [ˈkaØa], which incurs an additional penalty of 5.44.

Because the maxent learner's job is to maximize the probability of each candidate, it is able to make predictions about the occurrences of allophones [z] and [s^{θ}] despite their being harmonically bound.[35] While these candidates are never the most harmonic even in a probabilistic grammar, the maxent algorithm allows us to see where they are most likely to occur when taking only linguistic factors into consideration. While the base grammar is able to account for occurrences of [z] and [s^{θ}] to some degree, these allophones are particularly sensitive to social variation and will be reexamined in depth in Sections 6.5 and 7.3.

6.4.3.2 Configurations favoring debuccalization and deletion

Now that we have examined constraint configurations in which some degree of oral constriction is preferred, let's turn our attention to those that favor debuccalization and/or deletion of /s/. We begin by examining Table 24 and Table 25 below, which show tableaux for *buscan*

---

[35] An output candidate is said to be harmonically bound when no configuration of pertinent constraints will render it the winner of a given tableau.

/buskan/ ('they/y'all look for') and *aspecto* /aspekto/ ('aspect'), the two environments most favorable for the voiceless glottal fricative, [h].

Table 24

Tableau for /buskan/ (‘look for’; 3rd person plural)

/buskan/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesure)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.74	0.52	5.14	0.30	0.00	3.47	0.42	0.84	2.82	1.57			
a. ˈbus.kan	*	*	*	*	*						9.70	0.20	0
b. ˈbuz.kan	*		*	*		*			*		15.47	0.00	0
c. ˈbus$^{\theta}$.kan	*	*	*			*	*				13.29	0.01	0
d. ˈbuh.kan	*	*				*	*	*			8.99	0.40	5
e. ˈbuɦ.kan	*					*	*	*	*		11.29	0.04	1
f. ˈbuØ.kan						*	*	*	*	*	9.12	0.35	4

Table 25

Tableau for /aspekto/ (‘aspect’)

/aspekto/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesure)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.74	0.72	4.05	0.86	0.86	3.49	0.80	0.46	3.02	1.17			
a. asˈpek.to	*	*	*	*	*						10.23	0.13	3
b. azˈpek.to	*		*	*		*			*		15.16	0.00	0
c. as$^{\theta}$ˈpek.to	*	*	*			*	*				12.80	0.01	0
d. ahˈpek.to	*	*				*	*	*			9.21	0.36	3
e. aɦˈpek.to	*					*	*	*	*		11.51	0.04	0
f. aØˈpek.to						*	*	*	*	*	8.94	0.47	5

Table 24 and Table 25 are illustrative of two configurations in which debuccalization and deletion are the most probable outcomes, although [s] is still somewhat likely to occur. Both of these tableaux show /s/ occurring syllable-finally, which is a more salient than word-final position but significantly less salient than phrase-, word-, or syllable-initial position, rendering [h], [∅], and [s] the three most probable outputs. When we compare the predicted probabilities of these three allophones in Table 24 and Table 25, we can make two important observations. First, because the effort-cost of [s] is higher in Table 25 due to the height of the preceding vowel, its probability drops 7%. Second, because the weight of PRESERVE(segmental) is lower in the less salient, atonic prosodic position, [∅] is rendered more harmonic than [h].

Furthermore, while debuccalization is quite probable in the above examples, [ɦ] remains an unlikely candidate in both tableaux. While it always does better than [h] with respect to effort-cost, the penalty for LAZY(glottal_abduction) is too low preceding a voiceless consonant to make up for the penalty incurred by [ɦ] for PRESERVE(voice), which is weighted fairly high in syllable-final position. In these tableaux, we see that the learner is likely to choose the easiest candidate, the most salient candidate, or a happy medium in which effort-cost is significantly reduced by avoiding the high penalty for LAZY(oral_gesture) while still preserving some input specifications.

Let's now turn our attention to configurations that give highest preference to the two weakest allophones, [ɦ] and [∅]. Because the need to alleviate the effort cost of glottal abduction is highest before a voiced consonant, these contexts prove to be the most amenable for the voiced glottal fricative, which is assigned a probability of just over 25% in these environments. Table 26 displays the tableau for *dos* /dos/ ('two') when it occurs before a consonant that is [-coronal] and [+voice]. As the reader can see, the high penalty assigned to LAZY(glottal_abduction) in this context renders [ɦ] far more likely than any other output aside from [∅], which is the most harmonic candidate because LAZY(any_gesture) outranks PRESERVE(segmental) in this prosodic position. Furthermore, Table 27 serves to reinforce the importance of avoiding glottal abduction before a voiced consonant, as the predictions look quite different if /dos/ is followed by a consonant that is [-coronal] and [-voice]. In this tableau, [ɦ] is still a good candidate with a predicted probability of 17%; however, without the hefty penalty for LAZY(glottal_abduction), more-faithful [h] becomes slightly more harmonic than its voiced counterpart.

Table 26

Tableau for /dos/ ('two') when /s/ is followed by a consonant that is [-coronal] and [+voice]

/dos# b, m, g/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesure)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (voice, strid)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)			
weight	3.74	3.59	2.14	3.52	0.00	2.93	1.00	0.95	1.14	2.79	Harmony	Predicted	Observed
a. dos	*	*	*	*	*						12.99	0.01	1
b. doz	*		*	*		*			*		13.47	0.01	0
c. dos$^{\theta}$	*	*	*			*	*				13.40	0.01	0
d. doh	*	*				*	*	*			12.21	0.02	1
e. doɦ	*					*	*	*	*		9.76	0.26	3
f. doØ						*	*	*	*	*	8.81	0.69	6

Table 27

Tableau for /dos/ ('two') when /s/ is followed by a consonant that is [-coronal] and [-voice]

/dos# p, k/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesure)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (voice, strid)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)			
weight	3.74	0.72	4.05	0.86	0.86	2.93	1.00	0.95	1.14	2.79	Harmony	Predicted	Observed
a. dos	*	*	*	*	*						10.23	0.11	1
b. doz	*		*	*		*			*		12.72	0.01	0
c. dos$^{\theta}$	*	*	*			*	*				12.44	0.01	0
d. doh	*	*				*	*	*			9.34	0.26	0
e. doɦ	*					*	*	*	*		9.76	0.17	0
f. doØ						*	*	*	*	*	8.81	0.44	3

Table 26 and Table 27 serve to demonstrate the prevalence of debuccalization and deletion in the least salient prosodic position, word-final. Furthermore, because low-ranking Faithfulness constraints have little power to block lenition in this position, predicted probabilities of [h], [ɦ] and [∅] depend primarily on the effort cost of the context in which the given /s/ occurs. Deletion will always be the very best choice in this unsalient prosodic position as PRESERVE(segment) will never outrank LAZY(any_gesture); however, as the effort cost of [h] and [ɦ] changes according to context, deletion becomes more or less likely in turn.

We now turn our attention to the effects of various intervocalic contexts on predicted probabilities for word-final /s/. As the reader will remember, Markedness constraint weights for intervocalic contexts were highest when both vowels were [-high], lowest when both vowels were [+high], and intermediate for contexts in which one vowel was [-high] and the other was [+high]. Table 28, Table 29, Table 30, and Table 31 display tableaux representing each of the four possible intervocalic contexts in tonic, word-final position: *diez* /djes# e, o, a/ ('ten'), *dios* /djos# i, u/ ('god'), *país* /pais# e, o, a/ ('country'), and *país* /pais# i, u/, respectively. As was the case with Table 26 and Table 27 above, Faithfulness constraints are weighted fairly low in this prosodic position and therefore variations in Markedness constraint weights play a major role in determining predicted probabilities.

Table 28

Tableau for /djes/ ('ten') when /s/ is followed by a vowel that is [-high]

/djes# e, o, a/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.74	1.66	2.08	1.86	0.34	2.93	1.00	0.95	1.14	2.79			
a. djes	*	*	*	*	*						9.68	0.19	1
b. djez	*		*	*		*			*		11.75	0.02	0
c. djes$^{\theta}$	*	*	*			*	*				11.41	0.03	0
d. djeh	*	*				*	*	*			10.28	0.11	1
e. djeɦ	*					*	*	*	*		9.76	0.18	3
f. djeØ						*	*	*	*	*	8.81	0.46	2

Table 29

Tableau for /djos/ ('god') when /s/ is followed by a vowel that is [+high]

/djos# i, u/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.74	1.75	1.41	1.88	0	2.93	1.00	0.95	1.14	2.79			
a. djos	*	*	*	*	*						8.78	0.36	0
b. djoz	*		*	*		*			*		11.1	0.04	0
c. djos$^{\theta}$	*	*	*			*	*				10.83	0.05	0
d. djoh	*	*				*	*	*			10.37	0.07	0
e. djoɦ	*					*	*	*	*		9.76	0.14	1
f. djoØ						*	*	*	*	*	8.81	0.35	0

Table 30

Tableau for /pais/ (‘country’) when /s/ is followed by a vowel that is [-high]

/pais# e, o, a/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.74	1.67	1.84	1.67	0.00	2.93	1.00	0.95	1.14	2.79	Harmony	Predicted	Observed
a. paˈis	*	*	*	*	*						8.92	0.33	2
b. paˈiz	*		*	*		*			*		11.32	0.03	0
c. paˈis$^{\theta}$	*	*	*			*	*				11.18	0.03	0
d. paˈih	*	*				*	*	*			10.29	0.08	0
e. paˈiɦ	*					*	*	*	*		9.76	0.14	3
f. paˈiØ						*	*	*	*	*	8.81	0.37	3

Table 31

Tableau for /pais/ (‘country’) when /s/ is followed by a vowel that is [+high]

/pais# i, u/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.74	1.94	0.00	1.31	0.13	2.93	1.00	0.95	1.14	2.79	Harmony	Predicted	Observed
a. paˈis	*	*	*	*	*						7.12	0.66	1
b. paˈiz	*		*	*		*			*		9.12	0.09	0
c. paˈis$^{\theta}$	*	*	*			*	*				9.61	0.06	0
d. paˈih	*	*				*	*	*			10.56	0.02	0
e. paˈiɦ	*					*	*	*	*		9.76	0.04	0
f. paˈiØ						*	*	*	*	*	8.81	0.12	0

As the reader can see, [∅] is the most likely candidate in Table 28 in which /s/ is flanked by two vowels that are [-high]. In this context and prosodic position, predicted probabilities for [s] and [ɦ] are virtually equal, suggesting that the grammar considers the effort cost of producing [ɦ] to be approximately as bad as the loss of input [+strident], [+coronal], and [-voice] specifications in this unsalient context. However, as the vowel context becomes less effortful as seen in Table 29 and Table 30, [s] overtakes [ɦ] as the second-best candidate after [∅]; in these contexts, it appears that the perceptual cost of losing so many input specifications is only justified if the effort cost is eliminated entirely. By the last tableau, in which /s/ is flanked by two high vowels, [s] has overtaken [∅] as the most harmonic candidate. Here, even [z] and [s^{θ}] have overtaken the debuccalized allophones with respect to predicted probabilities, as the perceptual costs have well exceeded the effort costs.

6.4.3.3 Special cases

In this section, I briefly examine two special cases in the data that I argue are well accounted for by the maxent grammar. The first case is the /s(#)t/ context, which is particularly lenition-resistant compared to other preconsonantal contexts. Table 32 and Table 33 show representative tableaux in which /s/ occurs in two distinct prosodic positions: *está* /esta/ (‘to be’; 3rd person singular) in atonic, syllable-final position and *los* /los/ (‘the’; masculine, plural) in atonic, word-final position.

Table 32

Tableau for /esta/ ('to be'; 3rd person singular)

/esta/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (voice, strid)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)			
weight	3.74	1.04	2.46	0.00	0.00	3.49	0.80	0.46	3.02	1.17	Harmony	Predicted	Observed
a. esˈta	*	*	*	*	*						7.24	0.76	157
b. ezˈta	*		*	*		*			*		12.71	0.00	0
c. es$^{\theta}$ˈta	*	*	*			*	*				11.53	0.01	1
d. ehˈta	*	*				*	*	*			9.53	0.08	10
e. eɦˈta	*					*	*	*	*		11.51	0.01	2
f. eØˈta						*	*	*	*	*	8.94	0.14	32

Table 33

Tableau for /los/ ('the'; masculine, plural) when /s/ is followed by /t/

/los# t/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (voice, strid)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)			
weight	3.74	1.04	2.46	0.00	0.00	2.89	1.39	0.00	1.29	2.42	Harmony	Predicted	Observed
a. los	*	*	*	*	*						7.24	0.55	19
b. loz	*		*	*		*			*		10.38	0.02	1
c. los$^{\theta}$	*	*	*			*	*				11.52	0.01	0
d. loh	*	*				*	*	*			9.06	0.09	1
e. loɦ	*					*	*	*	*		9.31	0.07	1
f. loØ						*	*	*	*	*	7.99	0.26	2

While predicted probabilities differ in the above tableaux according to the strength of the prosodic position in which /s/ occurs, probabilities for output [s] are strikingly high as compared even to most prevocalic contexts. The effort cost of the /s(#)t/ sequence is so low that even in the least salient prosodic position, exemplified in Table 33, the voiceless sibilant has more than a 50% predicted probability. These patterns speak to a special articulatory nature of this sequence, which will be discussed in depth in Chapter 7.

The second special case is phrase-final /s/, which is unique in that it occurs in a weak prosodic position (word-finally) but at a phrase boundary, which is often marked by phrase-final lengthening and considered more salient than other word-final environments. As I discuss in Chapter 2, studies on Spanish /s/ lenition have produced inconsistent results with respect to phrase-final /s/ weakening, with some reporting near-categorical elision and others reporting near-categorical retention of [s]. Therefore, I believe it's valuable to briefly discuss the predictions made by the maxent grammar with respect to this controversial prosodic position. Table 34, Table 35, Table 36, and Table 37 show tableaux of four words in which /s/ occurs in phrase-final position. Because constraint weight exploration in Section 6.4.2 reveled an important effect of tonicity in syllable-final position, Table 34 and Table 36 show predicted probabilities when /s/ occurs in a tonic syllable, which can be compared to Table 35 and Table 37, in which /s/ occurs in an atonic syllable. For illustrative purposes, the first two tableaux examine the $V_{[-high]}$ __## context, while the latter two examine the $V_{[+high]}$ __##context.

Table 34

Tableau for /ingles/ (‘English’) in phrase-final position

/ingles##/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (voice, strid)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)			
weight	3.74	0.00	2.05	0.00	0.34	4.05	0.24	0.27	1.37	1.33	Harmony	Predicted	Observed
a. iŋˈles	*	*	*	*	*						6.13	0.67	10
b. iŋˈlez	*		*	*		*			*		11.21	0.00	0
c. iŋˈles$^{\theta}$	*	*	*			*	*				10.08	0.01	1
d. iŋˈleh	*	*				*	*	*			8.3	0.08	1
e. iŋˈleɦ	*					*	*	*	*		9.67	0.02	0
f. iŋˈleØ						*	*	*	*	*	7.26	0.22	7

Table 35

Tableau for /aɲos/ (‘years’) in phrase-final position

/aɲos##/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (voice, strid)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)			
weight	3.74	0.00	2.05	0.00	0.34	3.25	1.60	0.00	1.24	0.00	Harmony	Predicted	Observed
a. aˈɲos	*	*	*	*	*						6.13	0.46	7
b. aˈɲoz	*		*	*		*			*		10.28	0.01	1
c. aˈɲos$^{\theta}$	*	*	*			*	*				10.64	0.01	0
d. aˈɲoh	*	*				*	*	*			8.59	0.04	1
e. aˈɲoɦ	*					*	*	*	*		9.83	0.01	0
f. aˈɲoØ						*	*	*	*	*	6.09	0.48	13

Table 36

Tableau for /pais/ (‘country’) in phrase-final position

/pais##/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (voice, strid)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)	Harmony	Predicted	Observed
weight	3.74	1.83	0.00	0.35	0.00	4.05	0.24	0.27	1.37	1.33			
a. paˈis	*	*	*	*	*						5.92	0.74	28
b. paˈiz	*		*	*		*			*		9.51	0.02	1
c. paˈis$^{\theta}$	*	*	*			*	*				9.86	0.01	1
d. paˈih	*	*				*	*	*			10.13	0.01	0
e. paˈiɦ	*					*	*	*	*		9.67	0.02	1
f. paˈi∅						*	*	*	*	*	7.26	0.19	8

Table 37

Tableau for /ɡratis/ (‘free’) in phrase-final position

/ɡratis##/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (voice, strid)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)	Harmony	Predicted	Observed
weight	3.74	1.83	0.00	0.35	0.00	3.25	1.60	0.00	1.24	0.00			
a. ˈɡra.tis	*	*	*	*	*						5.92	0.51	2
b. ˈɡra.tiz	*		*	*		*			*		8.58	0.04	0
c. ˈɡra.tis$^{\theta}$	*	*	*			*	*				10.42	0.01	0
d. ˈɡra.tih	*	*				*	*	*			10.42	0.01	0
e. ˈɡra.tiɦ	*					*	*	*	*		9.83	0.01	0
f. ˈɡra.ti∅						*	*	*	*	*	6.09	0.43	1

In the tableaux above, the reader can see that probabilities are essentially split between [s] and [Ø], reflecting previous work on Spanish /s/ lenition in phrase-final position. More so than in any other position, it appears that intermediate allophones are highly improbable here, suggesting that the context in which /s/ occurs determines whether it will pattern like an /s/ at a phrase boundary (as in Table 34 and Table 36) or more like a word-final /s/, although predictions even in Table 35 and Table 37 still favor output [s] much more than other word-final examples we have examined. These findings will be explored in more detail in Chapter 7.

6.4.4 Assessing the base grammar's predictions

In order to identify areas where the grammar captures patterns in the data well and areas where it does less well, I compared predicted vs. observed probabilities for every possible context + prosodic position combination (n = 72). I anticipated that the correlation between predicted and observed probabilities would be fairly high because predicted probabilities are calculated based on constraint weights, which are determined based on observed data. However, examining the cases in which these values are most disparate reveals allophones, contexts, and prosodic positions in which the grammar had a more difficult time capturing important aspects of the observed data. The predicted and observed probabilities for all 72 context/position combinations by allophone are presented in Figure 57.

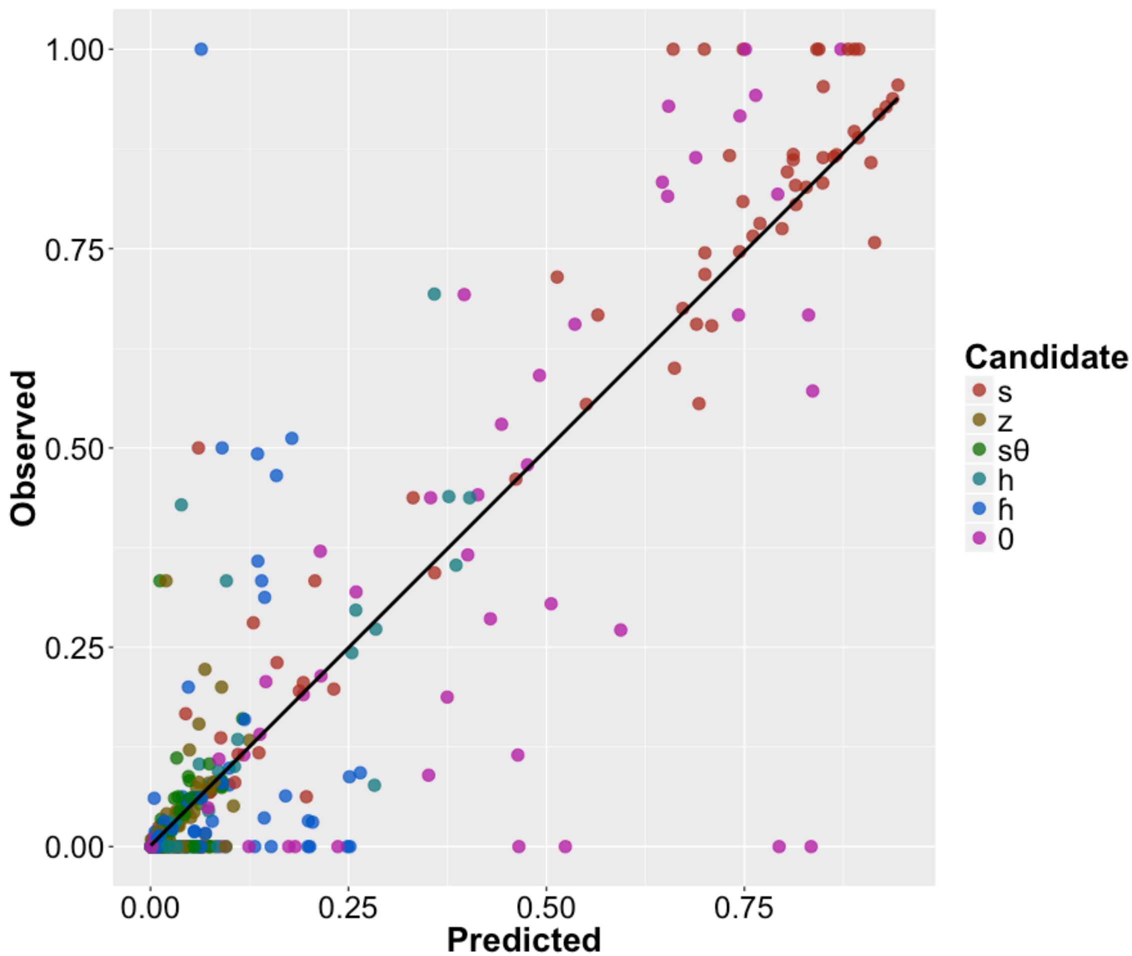

Figure 57. Predicted vs. observed probabilities for 72 context and prosodic position combinations; *r* = 0.91.

In Figure 57, the reader can see that the correlation between predicted and observed values is high (*r* = 0.91) but not perfect, which would indicate overfitting to the training data. Data points closest to the correlation line are the best matches in terms of predicted and observed values, and those furthest from the line in quadrants II and IV are the most disparate.

Figure 57 reveals that the grammar does the best at predicting probabilities for [s], [s^{θ}], and [z] and struggles the most at distinguishing between the two weakest allophones, [ɦ] and [Ø]. Specifically, the grammar frequently under-predicts [ɦ] in favor of [Ø], suggesting that—if these misalignments are occurring across the same context and/or prosodic position—some constraint differentiating between the two allophones may be missing from the grammar. In order to examine these mismatched data points in more depth, I separated the data first by context type and then by prosodic position, the results of which can be seen in Figure 58 and Figure 59.

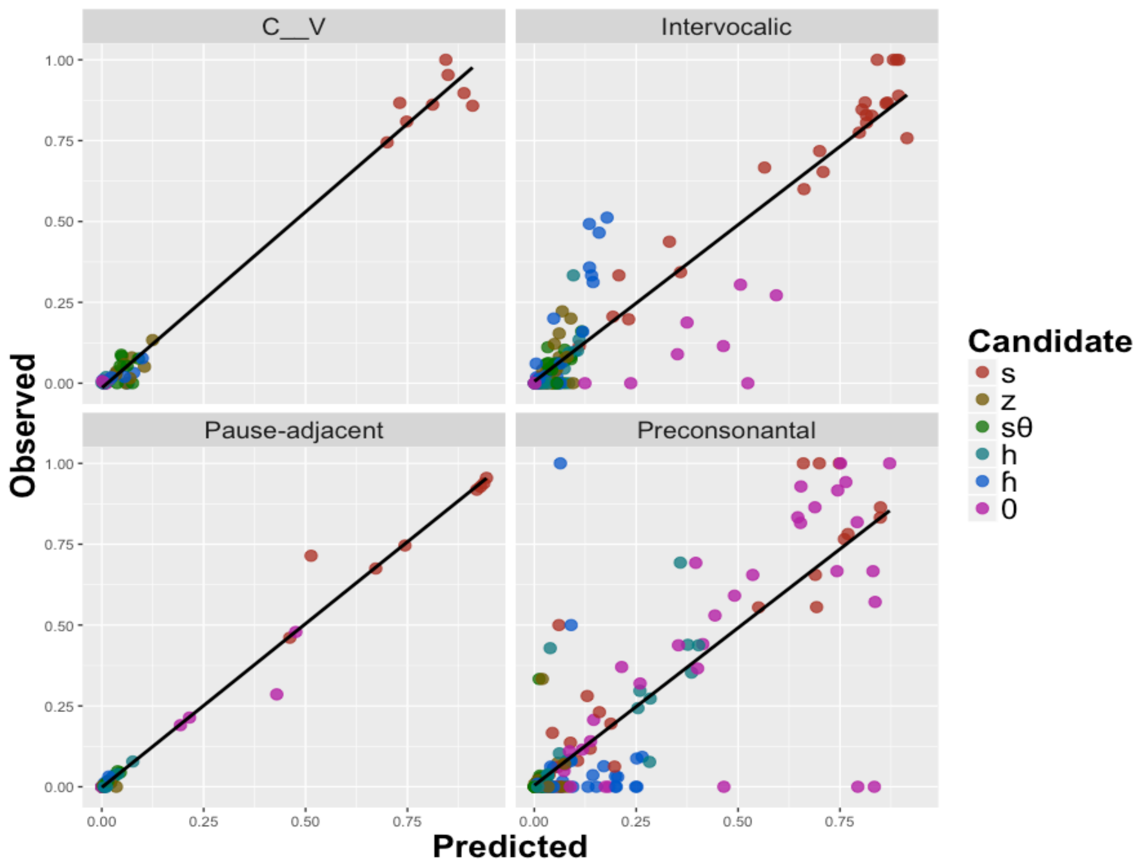

Figure 58. Predicted vs. observed probabilities by context type.

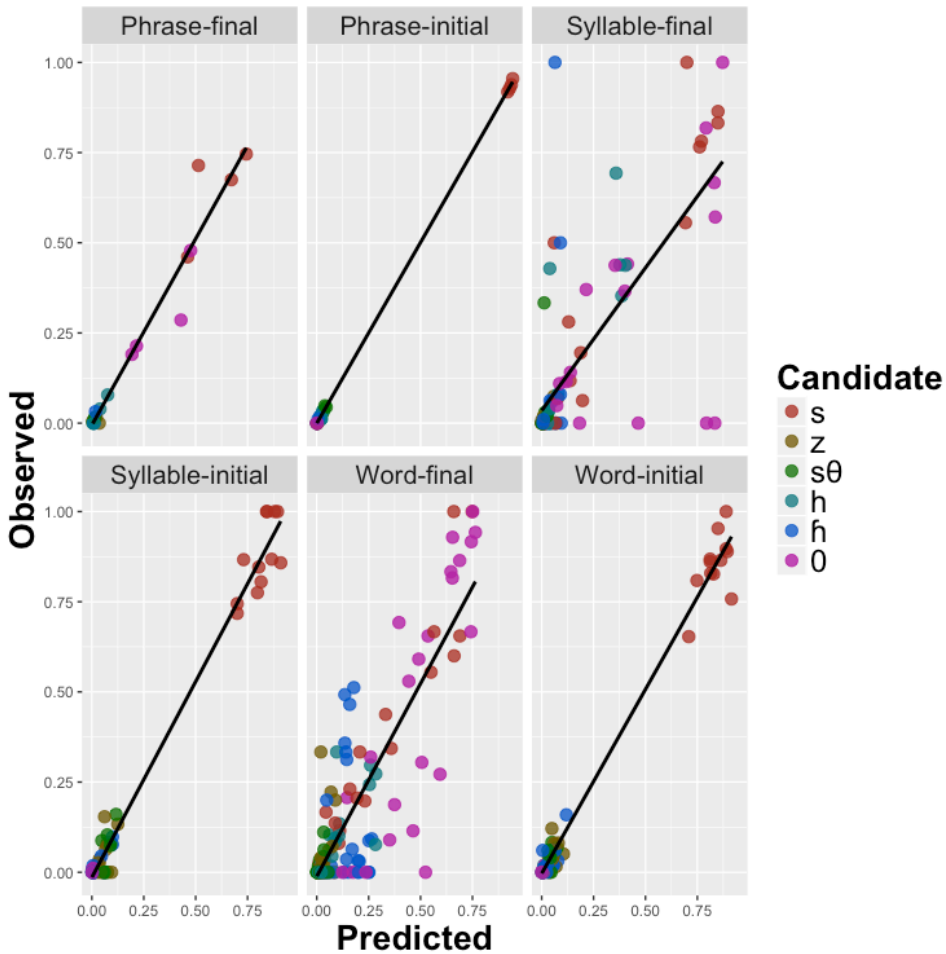

Figure 59. Predicted vs. observed probabilities by prosodic position.

Evident in Figure 58 and Figure 59 is that the model's over-prediction of [∅] occurs primarily in syllable-final position, which is always preconsonantal. Because [∅] is the most likely candidate in word-final, preconsonantal position, it appears that the grammar has overlearned preconsonantal deletion word-medially, despite the fact that it is much less prevalent in the data. Indeed, a comparison of the allophones that occur preconsonantally in syllable-final and word-final positions reveals that lenition in is much more wide-spread in the latter. Figure 60 visualizes this comparison.

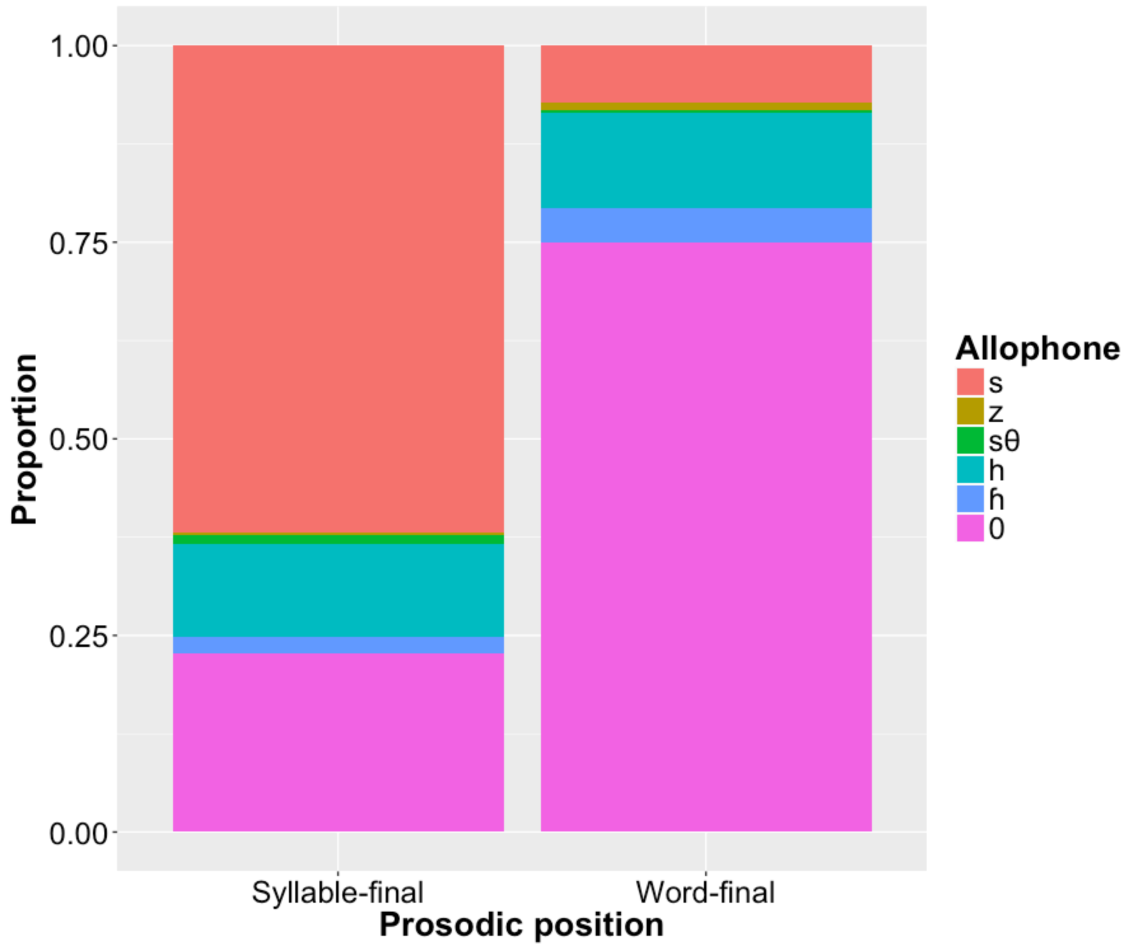

Figure 60. Comparison of allophones realized in preconsonantal, syllable- and word-final prosodic positions.

As the reader can see, syllable-final position—the more salient of the two—disfavors the weakest allophones [ɦ] and [∅] as compared to the less salient word-final position, even when phonological contexts are more or less controlled for. Therefore, because there are approximately 50% more tokens of preconsonantal, word-final /s/ (n = 2634) than tokens of syllable-final /s/ (n = 1696) in the observed data, it is possible that the grammar has assigned disproportionately high weights to preconsonantal Markedness constraints, rendering Faithfulness constraints in syllable-final position unable to outrank them at the rate needed to reflect patterns in the observed data. Aside from this over-prediction of deletion, however, predictions made by the grammar are a good reflection of the observed data.

The maxent grammar appears to be particularly adept at predicting use of the least-common allophones, [z] and [s^{θ}], which are both predicted and observed at probabilities of up to 25% in certain phonological contexts and prosodic positions. Even though the grammar makes predictions in such a way that [s] is always more likely to surface than these lenited allophones, the ability of the MaxEnt learner to account for cumulativity in the data allows it to still make good predictions about the contexts in which [z] and [s^{θ}] are most likely to surface, even if they do so at lower rates than [s].

Section 6.4 has presented various facets of the base grammar for Salvadoran /s/ lenition through a close examination of weights assigned to Markedness and Faithfulness constraints and the patterns we observe based on their interactions. In this dissertation, I argue that the most important features of the grammar are language-internal; that is, constraints that take only grammatical factors into consideration determine fundamental patterns of /s/ lenition, while the social characteristics of speakers then scale the importance of these constraints up or down. The following section examines changes in the relative importance of Faithfulness when constraints are modulated according to subgroups based on a speaker's region of origin, urban vs. rural place of origin, age, gender, and education.

6.5 Social factors

6.5.1 Region

The three regions examined in the present study were Santa Ana, San Salvador, and San Miguel. Hypotheses, which were based on Azcúnaga Lopez's (2010) linguistic atlas of El Salvador in tandem with language attitudes expressed by my participants during sociolinguistic interviews, stated that speakers from San Salvador would reflect the national dialectal standard, speakers from Santa Ana might show signs of contact with neighboring Guatemala, and speakers from San Miguel would lenite at the highest rates and use more non-standard variants in more salient prosodic positions.

Let's first examine general trends by looking at overall shifts the importance of Faithfulness in the grammar for the three regional groups. Table 38 below compares differences between Faithfulness constraint weights and Markedness constraint weights for each group as compared to the base grammar. This metric tells us the relative importance of preserving perceptual distinctions in input and output forms for each group; the *scaling factor* indicates the change in relative importance as compared to the Salvadoran speech community as a whole. Here, I remind the reader that, as described in Section 5.6, average Markedness and Faithfulness constraint weights were obtained via individual grammars fitted to each group.

Table 38

Relative importance of Faithfulness by region

Group	Average Markedness constraint weight	Average Faithfulness constraint weight	Relative importance of Faithfulness	Scaling factor
El Salvador (all)	1.29	2.16	0.87	N/A
Santa Ana	1.05	2.00	0.95	+0.08
San Miguel	1.41	2.13	0.72	-0.15
San Salvador	1.21	2.00	0.79	-0.08

Consistent with initial hypotheses, the relative importance of Faithfulness appears to be highest in the Santa Ana group, lowest in the San Miguel group, and about average (almost identical to the Salvadoran speech community as a whole) in the San Salvador group.

Next, let's turn our attention to changes in the relative importance of Faithfulness for individual constraints, the results of which present a more nuanced picture of regionally-conditioned /s/ lenition in El Salvador. Figure 61 below shows the scaling factor, or the change in relative importance, of the five Faithfulness constraints in the grammar by regional group.

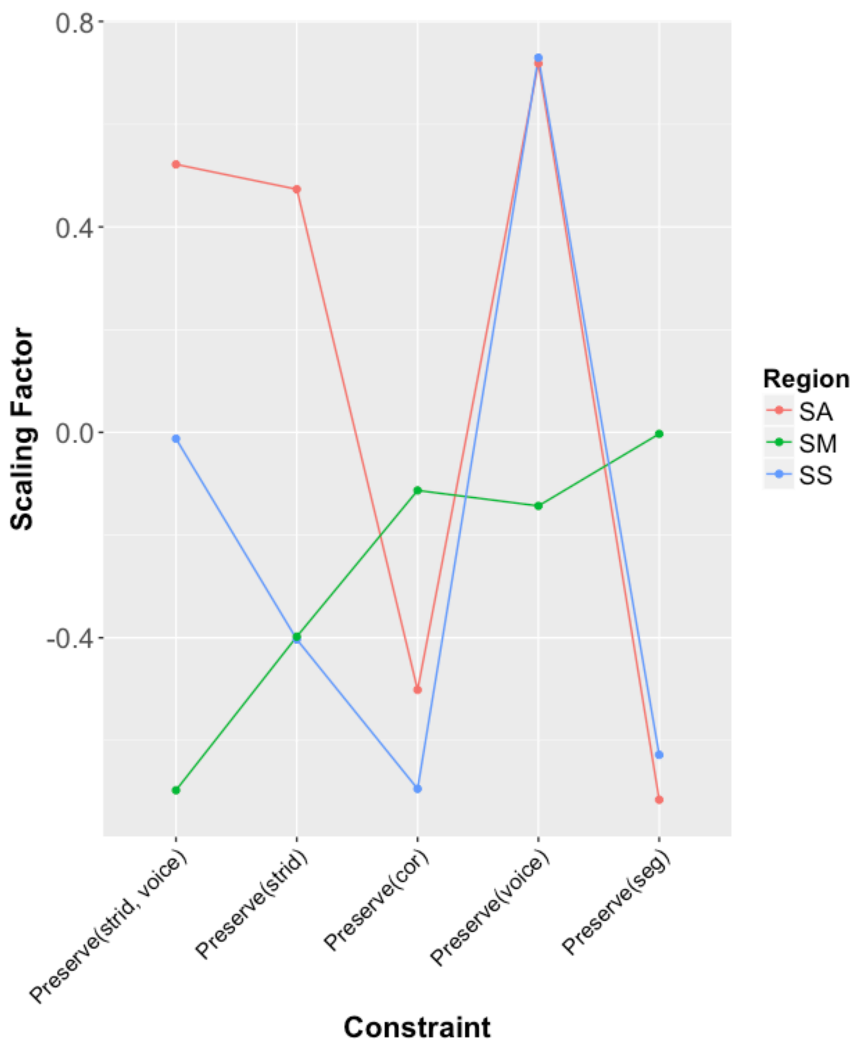

Figure 61. Changes in Faithfulness constraints compared to the base grammar for speakers from Santa Ana, San Miguel, and San Salvador.

In this figure, as with the previous table, a positive scaling factor indicates an increased importance of a constraint as compared to the base grammar; a negative scaling factor, of course, indicates the opposite. Strikingly, although speakers from Santa Ana seem to prioritize Faithfulness quite a bit more than speakers from San Salvador when it comes to preserving both features [+strident] and [-voice] or just the feature [+strident], the two groups pattern almost identically otherwise. Speakers from San Miguel, on the other hand, could not be more disparate: they not only prioritize Faithfulness less than the speech community as a whole across the board, but also do better on Faithfulness than the other two groups with respect to two interesting constraints: PRES(coronal), which serves to differentiate [s^{θ}] from [h], and PRESERVE(segmental), which differentiates [ɦ] from [∅]. Indeed, Figure 62 corroborates these observations in the raw data.

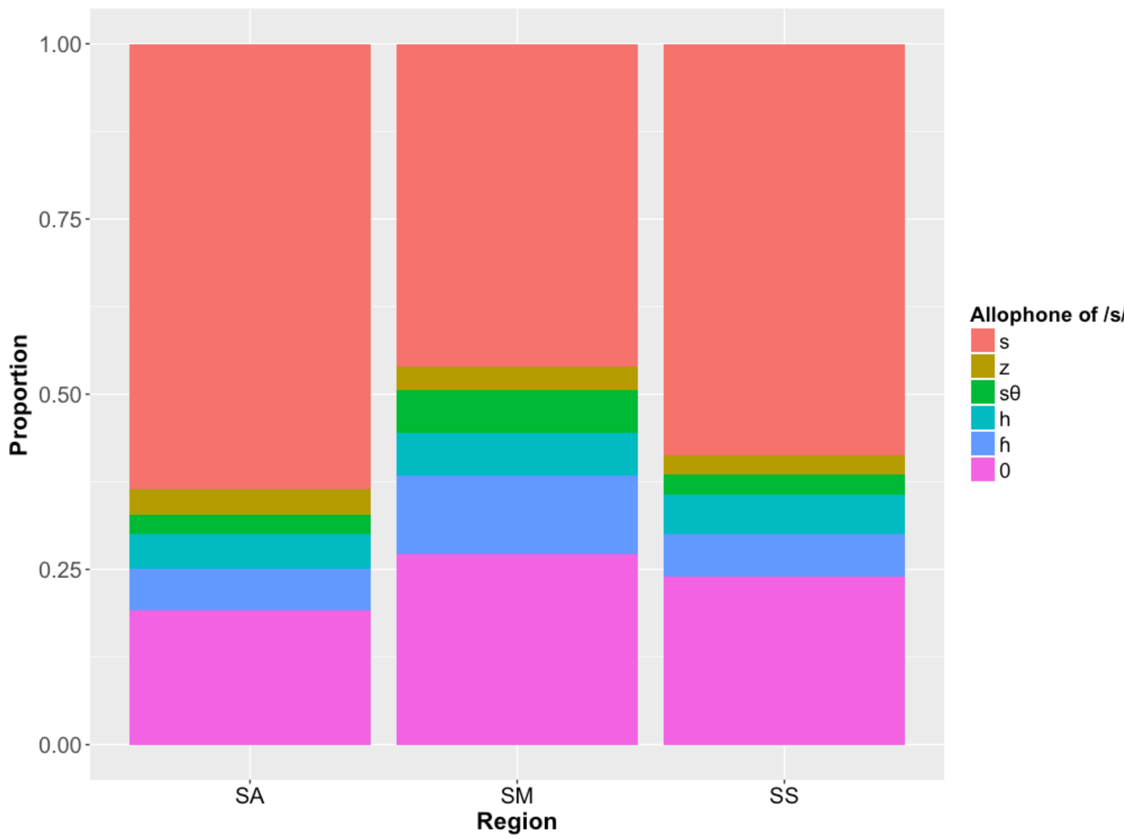

Figure 62. Proportion of allophones produced by region.

Here, we see that speakers from Santa Ana and San Salvador are akin in their production, although the latter produces slightly fewer allophones of [s] in favor of [∅]. Contrastingly, and as was reflected in the comparison of importance of Faithfulness constraints, Speakers from San Miguel produce a much wider range of intermediate allophones; in particular, [s^{θ}] and [ɦ].

These regional differences can also be seen in the predictions made by each scaled grammar. As hypothesized, speakers from San Miguel are more likely to weaken /s/ in stronger prosodic positions. Whereas the highest probability of producing [s^{θ}] in the base grammar is 12%, that number increases to 18% for speakers from San Miguel and decreases to just 8% in Santa Ana. More notable still is that, while the most likely prosodic position for this allophone in both regions is syllable-initial, it is most probable in tonic position in San Miguel and in atonic position in Santa Ana, furthering the argument that speakers from the former lenite in more prominent prosodic positions. Table 39 and Table 40 compare tableaux for words that favor the voiceless approximant in each region: *casa* /kasa/ ('house') in Santa Ana, and *hacer* /aseɾ/ ('to do/make') in San Miguel.

Table 39

Tableau for /kasa/ ('house'); Santa Ana

/kasa/		Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight		3.33	1.25	2.00	1.35	0.78	4.38	0.00	3.48	0.00	4.27			
a.	ˈka.sa	*	*	*	*	*						8.71	0.74	14
b.	ˈka.za	*		*	*		*			*		11.06	0.07	0
c.	ˈka.s$^{\theta}$a	*	*	*			*	*				10.96	0.08	9
d.	ˈka.ha	*	*				*	*	*			12.44	0.02	0
e.	ˈka.ɦa	*					*	*	*	*		11.19	0.06	1
f.	ˈka.∅a						*	*	*	*	*	12.13	0.02	0

Table 40

Tableau for /aseɾ/ ('to do/make'); San Miguel

/aseɾ/		Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight		3.82	1.67	1.86	2.67	0.30	4.32	0.00	4.49	0.70	4.38			
a.	aˈseɾ	*	*	*	*	*						10.32	0.71	41
b.	aˈzeɾ	*		*	*		*			*		13.37	0.03	1
c.	aˈs$^{\theta}$eɾ	*	*	*			*	*				11.67	0.18	19
d.	aˈheɾ	*	*				*	*	*			14.30	0.01	0
e.	aˈɦeɾ	*					*	*	*	*		13.33	0.04	1
f.	aˈ∅eɾ						*	*	*	*	*	13.89	0.02	0

Because the relative importance of Faithfulness is greater in Table 39 than in Table 40, output [ˈka.sa] in the Santa Ana grammar is more harmonic than [aˈseɾ] in the San Miguel grammar, yielding a greater probability that it will be produced. Furthermore, while predicted probabilities for unfaithful candidates are distributed fairly evenly in Table 39, the high ranking of PRESERVE(coronal) in Table 40 gives [aˈs^{θ}eɾ] the lion's share of the probability among non-[s] candidates.

Furthermore, while lenition in syllable-onset position is thought to be almost non-existent in most dialects of Spanish, speakers from San Miguel are predicted to debuccalize [s] at rates of more than 20% in syllable-onset position and up to 14% word-initially when flanked by two vowels that are [-high]. Table 41 and Table 42 show the predicted probabilities for two such words in the scaled grammar from San Miguel: *casa* /kasa/ ('house') and *son* /son/ ('to be'; third person plural).

Table 41

Tableau for /kasa/ ('house'); San Miguel

/kasa/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.82	1.67	1.86	2.67	0.30	3.96	0.37	3.26	0.00	6.18			
a. ˈka.sa	*	*	*	*	*						10.32	0.55	7
b. ˈka.za	*		*	*		*			*		12.31	0.08	0
c. ˈka.s$^{\theta}$a	*	*	*			*	*				11.68	0.14	0
d. ˈka.ha	*	*				*	*	*			13.08	0.03	0
e. ˈka.ɦa	*					*	*	*	*		11.41	0.19	2
f. ˈka.Øa						*	*	*	*	*	13.77	0.02	0

Table 42

Tableau for /son/ ('be'; third person plural) when preceded by a [-high] vowel; San Miguel

/e, o, a #son/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (voice, strid)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)	Harmony	Predicted	Observed
weight	3.82	1.67	1.86	2.67	0.30	4.78	0.00	3.48	0.00	6.48			
a. son	*	*	*	*	*						10.32	0.69	13
b. zon	*		*	*		*			*		13.13	0.04	1
c. s$^{\theta}$on	*	*	*			*	*				12.13	0.11	4
d. hon	*	*				*	*	*			13.75	0.02	0
e. ɦon	*					*	*	*	*		12.08	0.12	4
f. Øon						*	*	*	*	*	14.74	0.01	0

I first encourage the reader to compare Table 41 with Table 39 from above, as the two tableaux predict very different outcomes for the same word, with the only difference being the region of origin of the speaker. In Santa Ana, candidate [ˈka.sa] is almost 20% more likely than in San Miguel, where it forfeits some of its probability to more lenited outputs [ˈka.ɦa] and [ˈka.s^{θ}a]. In fact, more than one in five outputs of /kasa/ in San Miguel are predicted to be either [ˈka.ɦa] or [ˈka.ha]. For the word /son/, shown in Table 42, we see that Faithfulness constraints are more powerful in this more prominent prosodic position; however, [ɦon] still remains a fairly probable outcome in this context at 12%.

This section has examined the differences in the relative importance of Faithfulness constraints in three distinct regions of El Salvador and how those differences are reflected in predicted probabilities for output candidates. Furthermore, I believe this examination provides support for the proposal that fundamental patterns of the grammar are determined by linguistic factors and then simply modulated, or scaled, by social factors. This can be seen, for example, in the fact that even speakers from San Miguel, who lenite /s/ at much higher frequencies than speakers from the other two regions, still adhere to a prosodic hierarchy in determining which positions are more or less acceptable for lenition (compare Table 41 and Table 42 above). A discussion of the lenition patterns observed by speakers from different regions will be presented in the next chapter.

6.5.2 Origin

Within each of the three regions, participants came from both rural and urban places of origin. Based on existing literature on the urban/rural divide in the Spanish-speaking world, reviewed in Chapter 2, I hypothesized that speakers from rural areas would lenite at higher rates than those from urban communities, use more non-standard variants such as [s^{θ}], and use more lenited variants in more salient prosodic positions.

We first examine general trends by looking at overall shifts the importance of Faithfulness in the grammar for the three regional groups. As with the previous section, Table 43 compares ratios of Faithfulness constraint weights to Markedness constraint weights for the two groups as compared to the base grammar. I remind the reader that the *scaling factor* indicates the change in relative importance as compared to the Salvadoran speech community as a whole.

Table 43

Relative importance of Faithfulness by origin

Group	Average Markedness constraint weight	Average Faithfulness constraint weight	Relative importance of Faithfulness	Scaling factor
El Salvador (all)	1.29	2.16	0.87	N/A
Rural	1.37	2.13	0.76	-0.11
Urban	1.15	2.04	0.89	+0.02

While differences are much smaller than we observed with respect to region, comparisons in Table 43 confirm initial hypotheses that the relative importance of Faithfulness is higher for urban speakers than for the general speech community, and lower for rural speakers.

Next, let's turn our attention to changes in the relative importance of Faithfulness for individual constraints, the results of which present a more nuanced picture of origin-conditioned /s/ lenition in El Salvador. Figure 63 shows the scaling factor of the five Faithfulness constraints in the grammar by rural vs. urban origin.

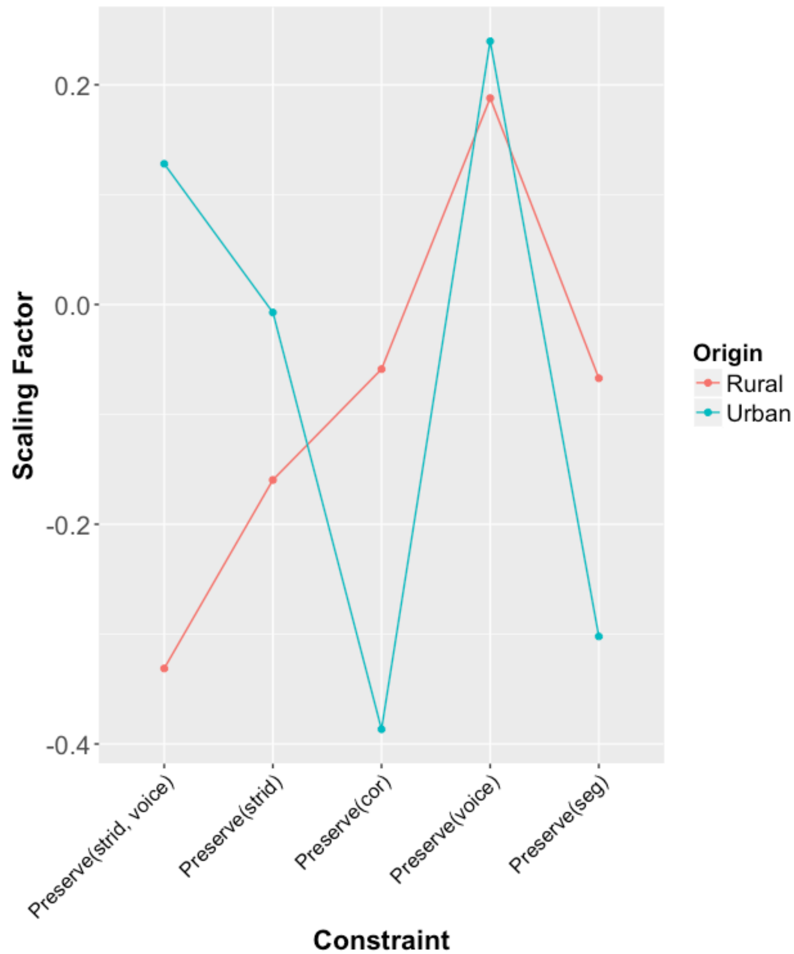

Figure 63. Changes in Faithfulness constraints compared to the base grammar for speakers from rural and urban origins.

Changes in Faithfulness constraints for rural and urban speakers parallel what we saw with respect to region in two important ways: here, as before, the group hypothesized to speak a more standardized dialect prioritizes preservation of the features [+strident] and [-voice], while the less standard group shows the greatest differences in the preservation of [+coronal] and [+segmental]. Again, this appears to be an accurate reflection of the raw data, which can be seen in Figure 64.

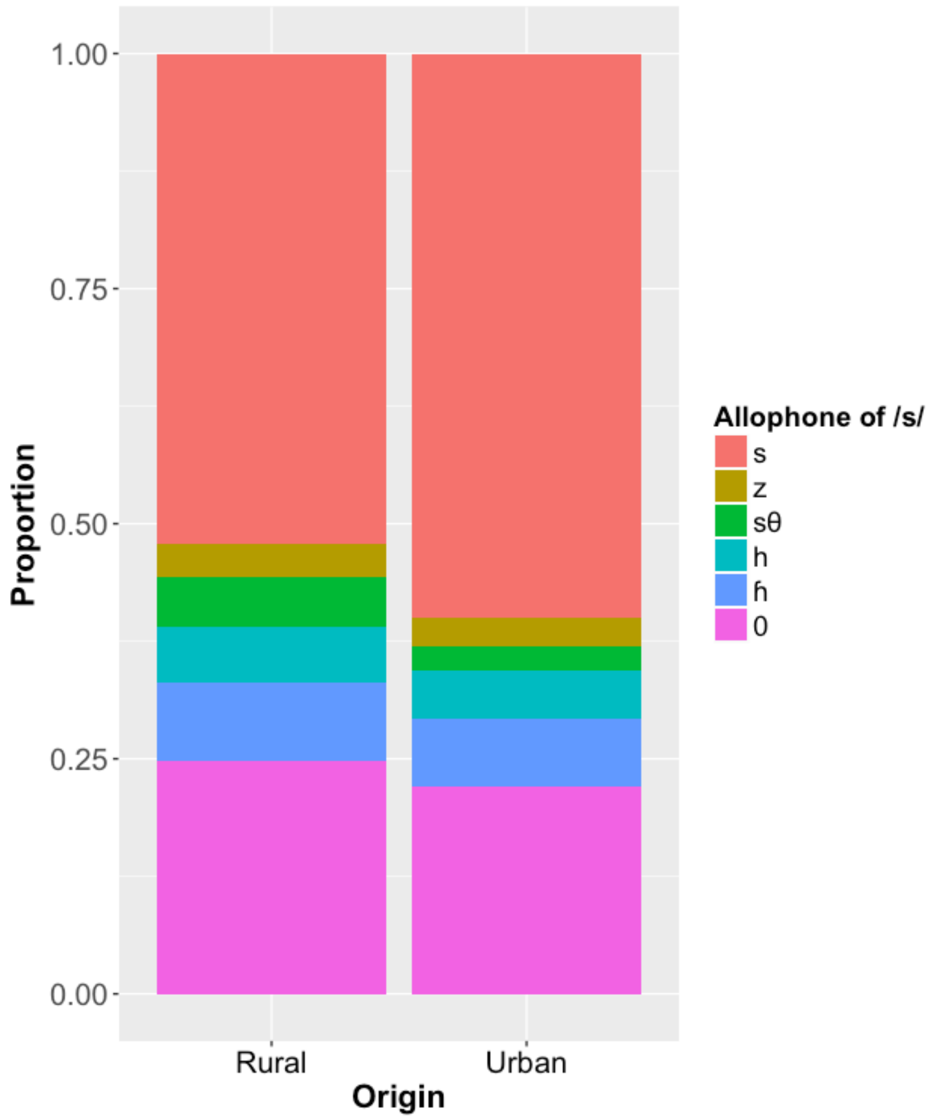

Figure 64. Proportion of allophones produced by origin.

Here we see that, like speakers from San Miguel, rural speakers produce much less [s] in general than their counterparts. While they also delete /s/ slightly more often, they produce a higher proportion of intermediate allophones, including [s^{θ}] and [ɦ], as the comparison of Faithfulness constraints in Figure 63 suggests.

These differences with respect to origin can also be seen in the predictions made by each scaled grammar. As hypothesized, rural speakers are more likely to weaken /s/ in stronger prosodic positions. Like San Migueleños, rural Salvadorans show a strong preference for [s^{θ}] as compared to their urban counterparts, with this allophone predicted at nearly 20% in syllable-initial, tonic position. In the same prosodic position, predicted probabilities drop to only 6% for urban speakers. Table 44 and Table 45 compare tableaux for the word *pasar* /pasaɾ/ ('to pass') for rural and urban speakers.

Table 44

Tableau for /pasaɾ/ (‘to pass’); rural

/pasaɾ/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (strid, voice)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)	Harmony	Predicted	Observed
weight	3.58	1.61	2.18	2.03	0.30	3.69	0.00	3.16	0.90	3.16			
a. paˈsaɾ	*	*	*	*	*						9.70	0.74	11
b. paˈzaɾ	*		*	*		*			*		12.38	0.05	0
c. paˈsᶿaɾ	*	*	*			*	*				11.06	0.19	2
d. paˈhaɾ	*	*				*	*	*			12.04	0.00	0
e. paˈɦaɾ	*					*	*	*	*		11.33	0.01	0
f. paˈØaɾ						*	*	*	*	*	10.91	0.01	0

Table 45

Tableau for /pasaɾ/ (‘to pass’); urban

/pasaɾ/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (strid, voice)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)	Harmony	Predicted	Observed
weight	3.59	1.55	1.85	1.75	0.35	4.83	0.00	3.85	0.71	3.85			
a. paˈsaɾ	*	*	*	*	*						9.09	0.88	12
b. paˈzaɾ	*		*	*		*			*		12.73	0.02	0
c. paˈsᶿaɾ	*	*	*			*	*				11.82	0.06	1
d. paˈhaɾ	*	*				*	*	*			13.82	0.01	0
e. paˈɦaɾ	*					*	*	*	*		12.98	0.02	0
f. paˈØaɾ						*	*	*	*	*	13.24	0.01	0

Upon comparing Table 44 and Table 45, we see that the significantly lower ranking of PRESERVE(strident, voice) in the former allows for [paˈs^{θ}aɾ] to become a much more harmonic candidate, and therefore a more probable output. The predicted probability of this candidate in Table 44, 19%, is even higher than in the San Miguel grammar. Additionally, the relatively low ranking of PRESERVE(strident, voice) in Table 44 as compared to Table 45 even results in an increased probability for output [paˈzaɾ] despite the higher ranking of PRESERVE(voice) in this grammar. As was the case with respect to region, the juxtaposition of these tableaux reveals two grammars that make very different predictions: while one predicts [s] almost categorically, the other predicts a lenited allophone at a rate of one in four.

Similarly, the rural grammar predicts much higher occurrence of lenition—whether it be gestural undershoot or debuccalization—in salient prosodic positions. Tonicity plays a particularly important role in determining the likelihood of debuccalized allophones in these positions; in some contexts, the predicted probability of producing [ɦ] in word-initial position doubles in atonic syllables. This is evidence that, while rural speakers do show a high proclivity to lenite /s/ in salient prosodic positions, they still adhere to conditioning factors in the base grammar, which determine just how probable these unconventional cases of lenition are. In order to illustrate this contrast, Table 46 and Table 47 compare rural predicted probabilities for *saben* /saben/ ('they/y'all know') and *saber* /saber/ ('to know') in identical, $V_{[-high]}$ __ $V_{[-high]}$ contexts.

Table 46

Tableau for /saben/ (‘know’; 3rd person plural) when /s/ is preceded by a vowel that is [-high]; rural

/e, o, a #saben/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (strid, voice)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.58	1.61	2.18	2.03	0.30	4.81	0.00	3.75	0.00	6.13	Harmony	Predicted	Observed
a. ˈsa.βen	*	*	*	*	*						9.70	0.80	1
b. ˈza.βen	*		*	*		*			*		12.60	0.04	0
c. ˈs$^{\theta}$a.βen	*	*	*			*	*				12.18	0.07	0
d. ˈha.βen	*	*				*	*	*			13.75	0.01	0
e. ˈɦa.βen	*					*	*	*	*		12.14	0.07	1
f. ˈØa.βen						*	*	*	*	*	14.69	0.01	0

Table 47

Tableau for /saberɾ/ (‘to know’) when /s/ is preceded by a vowel that is [-high]; rural

/e, o, a #saberɾ/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (strid, voice)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.58	1.61	2.18	2.03	0.30	4.04	0.13	3.44	0.00	5.67	Harmony	Predicted	Observed
a. saˈβeɾ	*	*	*	*	*						9.70	0.64	4
b. zaˈβeɾ	*		*	*		*			*		11.83	0.08	0
c. s$^{\theta}$aˈβeɾ	*	*	*			*	*				11.54	0.10	0
d. haˈβeɾ	*	*				*	*	*			12.80	0.03	1
e. ɦaˈβeɾ	*					*	*	*	*		11.19	0.14	1
f. ØaˈβeɾΓ						*	*	*	*	*	13.28	0.02	0

As the reader can see, the lower rankings of constraints PRESRVE(strident, voice), PRESERVE(coronal), and PRESERVE(segmental) in Table 47 allow for richer variation in predicted probabilities, with each candidate assigned a likelihood of greater than 1%. Furthermore, because PRESERVE(coronal) is ranked comparatively low in this atonic syllable, [ɦaˈβeɾ] is slightly more harmonic than [s^{θ}aˈβeɾ]. This is in contrast to what we see in Table 46, where a higher ranking of this constraint renders [ˈsθa.βen] and [ˈɦa.βen] equally probable. Additionally, the higher rankings of both PRESREVE(strident, voice) and PRESERVE(segmental) in this tableau make [ˈsa.βen] the best candidate by far, with likelihoods of candidates [ˈza.βen], [ˈha.βen], and [ˈØa.βen] halved as a result.

For input /e, o, a #saben/ produced by urban Salvadorans, the likelihood of output [ˈsa.βen] increases to 85%, and the stronger [ˈs^{θ}a.βen] overtakes [ˈɦa.βen] as the second-most harmonic candidate. For input /e, o, a #saber/, the likelihood of output [saˈβeɾ] increases to 79% and [ɦaˈβeɾ] drops to 8%. As this section has demonstrated, urban speakers in El Salvador prioritize Faithfulness significantly more than their rural counterparts. However, it is important to recognize that even this more 'standard' group is predicted to lenite /s/ in strong prosodic positions at rates that are much higher than what Lipski observed in the 1980s, or than is seen in most /s/ weakening dialects of Spanish.

6.5.3 Age

Within each rural or urban group, participants were balanced for age, with older participants 41 years of age or older at the time of participation and younger participants between the ages of 18 and 40. Based on existing literature on the relationship between age and /s/ lenition, it was difficult to make concrete hypotheses. While my own experiences in the community as well as sociocultural information obtained through sociolinguistic interviews suggested that younger speakers' production of /s/ was closer to a linguistic standard, the majority of work on age-graded sound change documents younger speakers innovating away from the prestigious variant (i.e., [s]). For a review of this literature, see Chapter 2.

As we did with both region and origin, let's first examine general trends in age by looking at overall shifts the importance of Faithfulness in the grammar. As with the previous sections, Table 48 below compares ratios of Faithfulness constraint weights to Markedness constraint weights for the two groups as compared to the base grammar.

Table 48

Relative importance of Faithfulness by age group

Group	Average Markedness constraint weight	Average Faithfulness constraint weight	Relative importance of Faithfulness	Scaling factor
El Salvador (all)	1.29	2.16	0.87	N/A
Older	1.32	2.04	0.72	-0.15
Younger	1.21	2.24	1.03	+0.16

Differences here reveal that that the relative importance of Faithfulness is higher for younger speakers than for the general speech community, and lower for older speakers. It is apparent that older speakers are the more linguistically non-standard group overall, and an examination of the relative importance of Faithfulness for individual constraints shows patterns that are very similar to those seen with respect to region and origin. In other words, older speakers pattern similarly to rural speakers and San Migueleños, whereas younger speakers pattern more similarly to speakers from San Salvador, Santa Ana, and urban communities. Figure 65 shows the scaling factor of the five Faithfulness constraints in the grammar by age group.

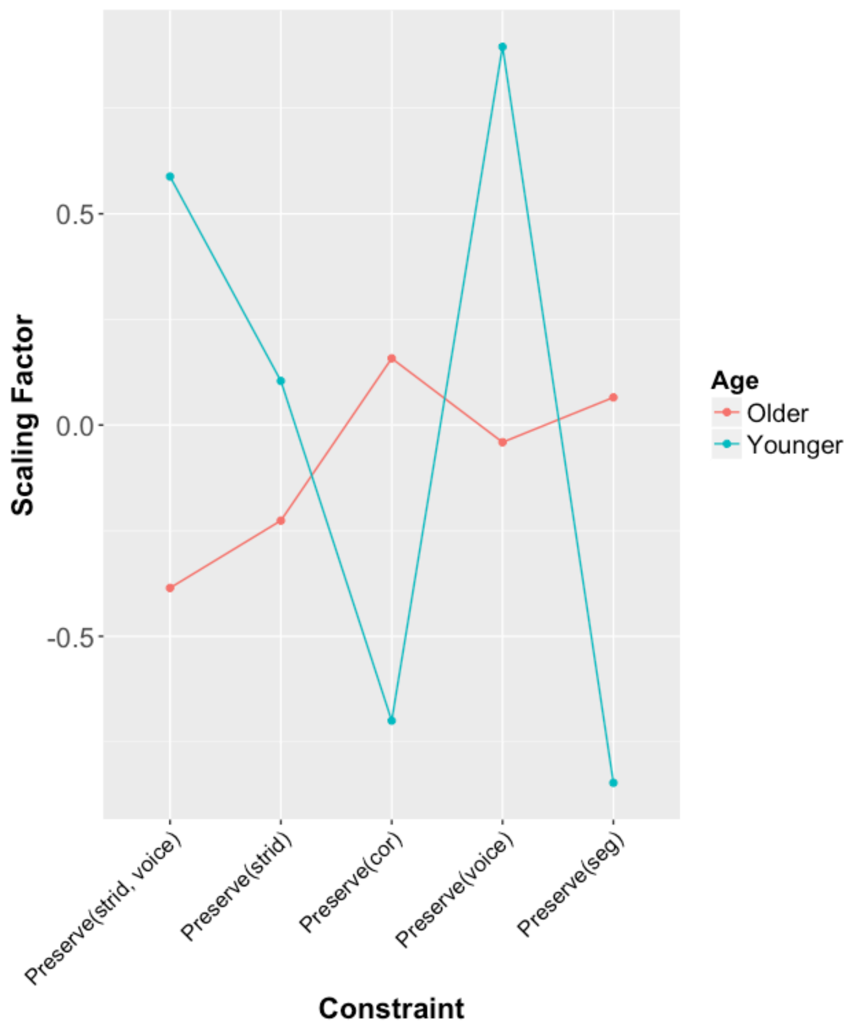

Figure 65. Changes in Faithfulness constraints compared to the base grammar for speakers from older and younger age groups.

Changes in Faithfulness constraints for older and younger speakers parallel what we saw with respect to region and origin in two important ways: here, as before, the group with more standardized speech prioritizes preservation of the features [+strident] and [-voice], while the less standard group shows the greatest differences in the preservation of [+coronal] and [+segmental]. Again, this appears to be an accurate reflection of the raw data, which can be seen in Figure 66.

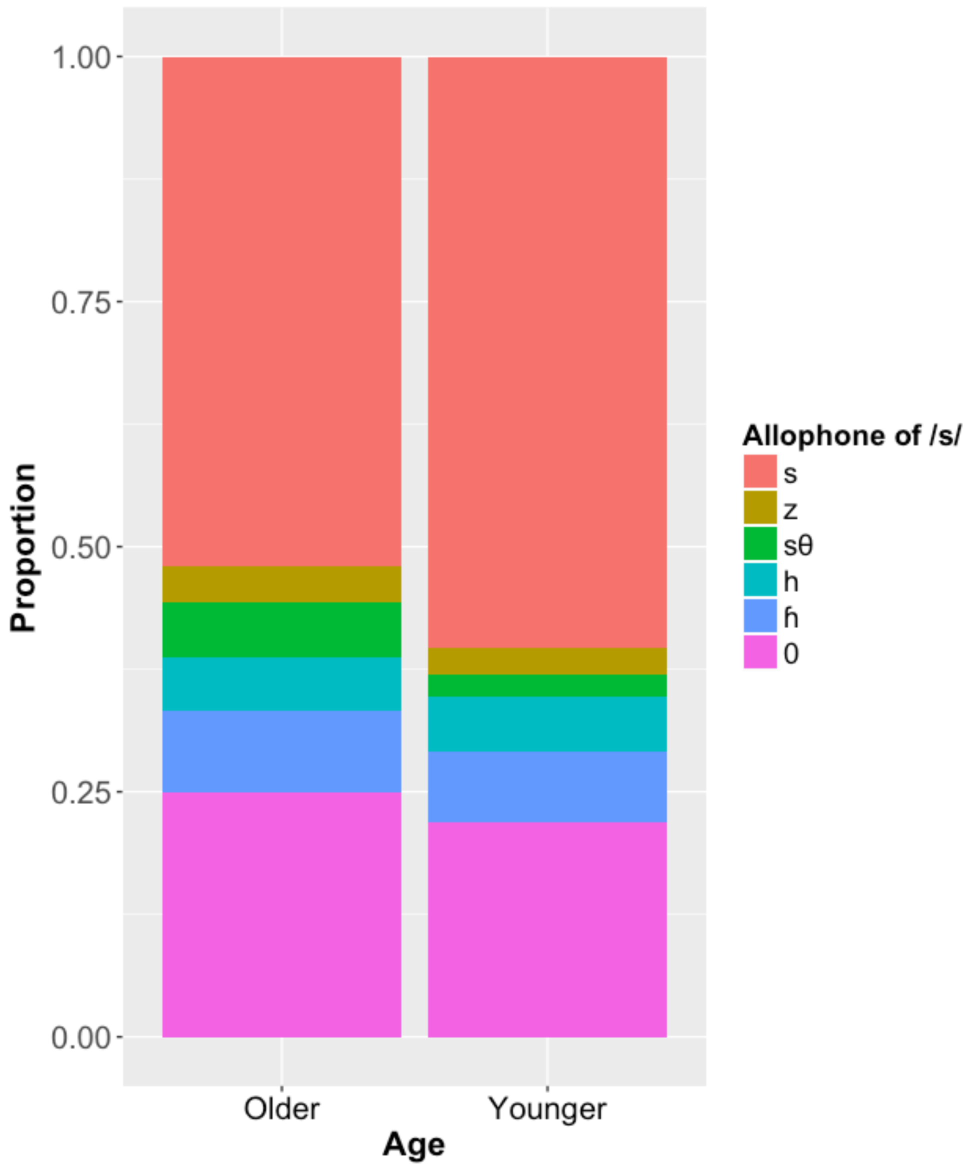

Figure 66. Proportion of allophones produced by age group.

Here, we see that older speakers utilize non-[s] allophones much more than younger speakers. In fact, the distribution of allophones in Figure 66 is almost identical to that of rural and urban speakers presented in Figure 64.

Akin to the other more innovative groups, older speakers in El Salvador lenite /s/ more frequently and more extremely in salient prosodic positions. In this section, we examine differences in predicted output probabilities for the strongest prosodic position: phrase-initial, tonic. While the base grammar assigns output [s] a likelihood of 93% in this position, we see these probabilities shift considerably when we examine older and younger speakers separately. Table 49 and Table 50 compare predicted probabilities for the inputs *cinco* /sinko/ (‘five’) and *cierto* /sjerto/ (‘sure/true’) in phrase-initial position for the two age groups.

Table 49

Tableau for /##sinko/ ('five'); older speakers

/##sinko/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (strid, voice)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.39	0.00	0.00	2.30	0.00	5.04	0.00	0.99	0.43	5.18	Harmony	Predicted	Observed
a. ˈsiŋ.ko	*	*	*	*	*						5.69	0.90	6
b. ˈziŋ.ko	*		*	*		*			*		11.16	0.00	0
c. ˈs$^{\theta}$iŋ.ko	*	*	*			*	*				8.43	0.06	0
d. ˈhiŋ.ko	*	*				*	*	*			9.42	0.02	0
e. ˈ ɦiŋ.ko	*					*	*	*	*		9.85	0.01	0
f. ˈØiŋ.ko						*	*	*	*	*	11.64	0.00	0

Table 50

Tableau for /##sjeɾto/ ('sure/true'); younger speakers

/##sjeɾto/	Lazy (any_ gesture)	Lazy (glot_ abduct)	Lazy (oral_ gesture)	Lazy (strid_ fric)	Lazy (vcl_ strid_fric)	Pres (strid, voice)	Pres (strid)	Pres (cor)	Pres (voice)	Pres (seg)			
weight	3.91	0.00	0.00	0.00	0.00	2.61	1.33	0.41	2.10	2.45	Harmony	Predicted	Observed
a. ˈsjeɾ.to	*	*	*	*	*						3.91	0.97	4
b. ˈsjeɾ.to	*		*	*		*			*		8.62	0.00	0
c. ˈs$^{\theta}$jeɾ.to	*	*	*			*	*				7.85	0.02	0
d. ˈhjeɾ.to	*	*				*	*	*			8.26	0.01	0
e. ˈɦjeɾ.to	*					*	*	*	*		10.36	0.00	0
f. ˈØjeɾ.to						*	*	*	*	*	8.90	0.00	0

Here, in this very prominent prosodic position, we see the differences in the relative importance of Faithfulness constraints at work. While Faithfulness is more than twice as important as Markedness in both Table 49 and Table 50, the distribution of Faithfulness penalties in the latter results in worse harmony scores for all non-[s] allophones, not just [∅], as is the case in former. While deletion here is technically more heavily penalized among older speakers, none of the other candidates suffer enormous additional penalties beyond violations incurred for PRESERVE(strident, voice), and a richer variation of outputs is therefore predicted by the grammar.

When we look at observed values for each group above, we see that both older and younger speakers produced [s] categorically in these words phrase-initially. However, even though older speakers never produced lenited allophones for this specific input, the learner has seen enough examples of phrase-initial weakening in other inputs to know that some variation is possible. For the younger speakers, however, so few examples of weakening in this salient position have been observed that the grammar predicts near-categorical production of [ˈsjeɾ.to]. Indeed, an examination of phrase-initial lenition in the raw data confirms that, while /s/ is rarely weakened, the majority of lenition cases pertain to speakers of the older age group. Figure 67 below shows these differences.

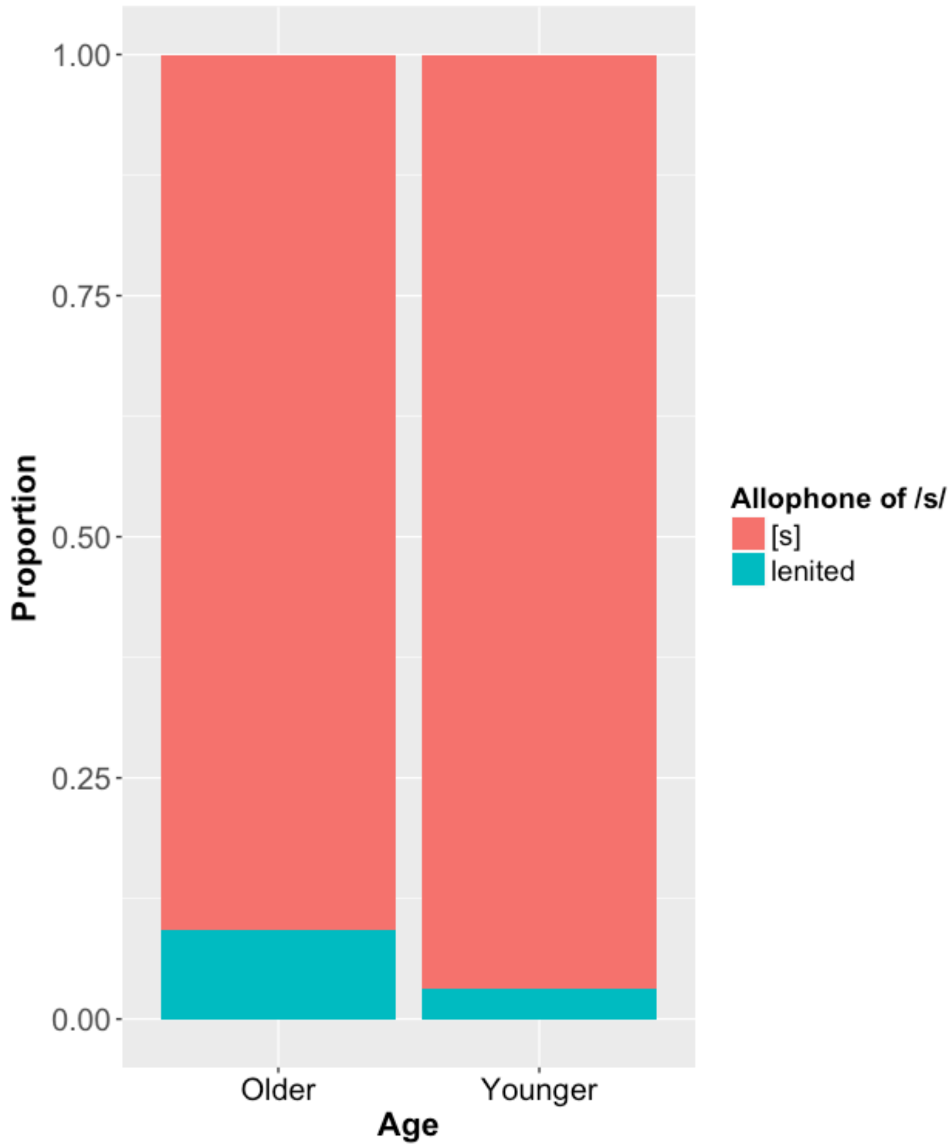

Figure 67. Proportion of lenited vs. non-lenited allophones produced in phrase-initial position by age group.

Here we see that the grammar does, in fact, do a good job of predicting future outputs of phrase-initial /s/ based on the observed data: older speakers lenite about 9% of the time, and younger speakers lenite 3% of the time, just as Table 49 and Table 50 predict.

6.5.4 Gender

Within each age group, participants were balanced for gender. Based on existing literature on the relationship between gender and phonological variation, I hypothesized that that men would lenite at higher rates and in more salient prosodic positions than women. Again, let's first examine general trends in gender differences by looking at overall shifts the importance of Faithfulness in the grammar. As with the previous sections, Table 51 compares ratios of Faithfulness constraint weights to Markedness constraint weights for the two genders as compared to the base grammar.

Table 51

Relative importance of Faithfulness by gender

Group	Average Markedness constraint weight	Average Faithfulness constraint weight	Relative importance of Faithfulness	Scaling factor
El Salvador (all)	1.29	2.16	0.87	N/A
Women	1.50	2.21	0.71	-0.16
Men	1.03	2.06	1.03	+0.16

Differences observed here reveal that Faithfulness is more important in the grammar of male Salvadorans, and that the difference between the two genders is the most distinct of any group comparison yet. While it is clear that women produce more non-standard variants of /s/ overall, an examination of the relative importance of Faithfulness for individual constraints, displayed in Figure 68, shows patterns that deviate slightly from those observed in previous sections.

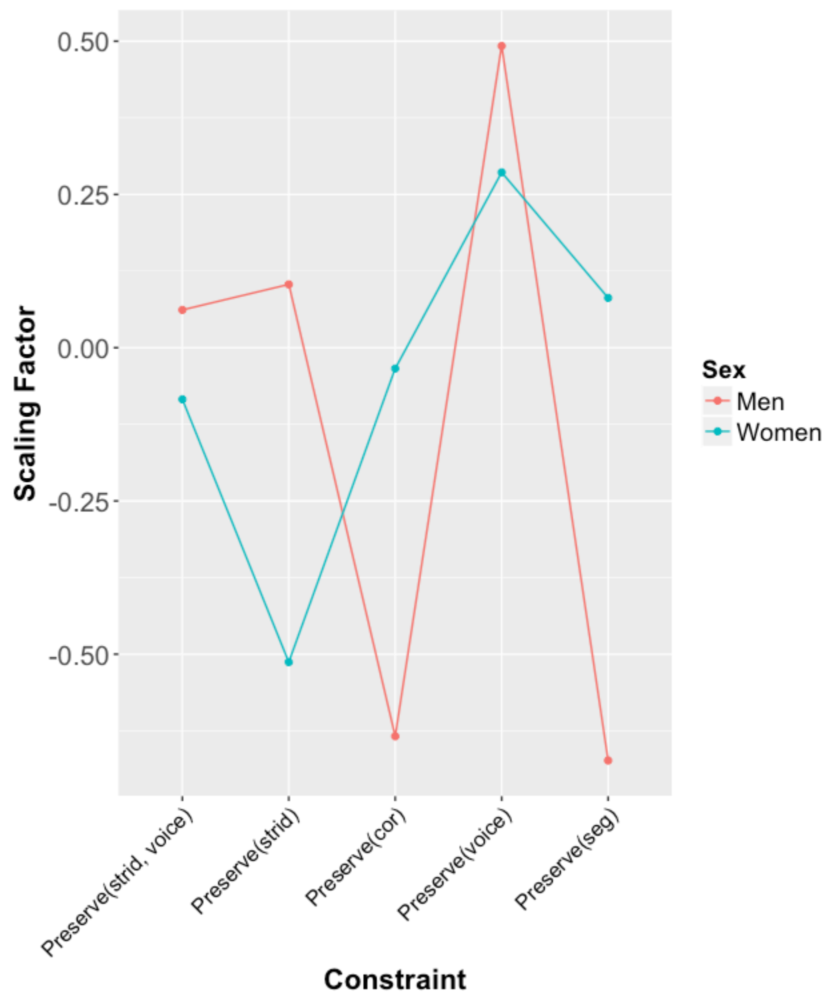

Figure 68. Changes in Faithfulness constraints compared to the base grammar for men and women.

Specifically, the difference in the relative importance of PRESERVE(strident) is much greater than for PRESERVE(strident, voice) in Figure 68, suggesting that the groups differ less in their production of [s] and more in their choice of intermediate allophones [z] vs. [s^θ]. This is the first social factor for which this is the case. If men are preserving the [+strident] input specification more than women and therefore producing more [z], it's curious that their scaling factor for PRESERVE(voice) is so large. This suggests that men likely prefer [h] over [ɦ] at a higher rate than women prefer [s^θ] over [z], an observation that is confirmed in the raw data in Figure 69.

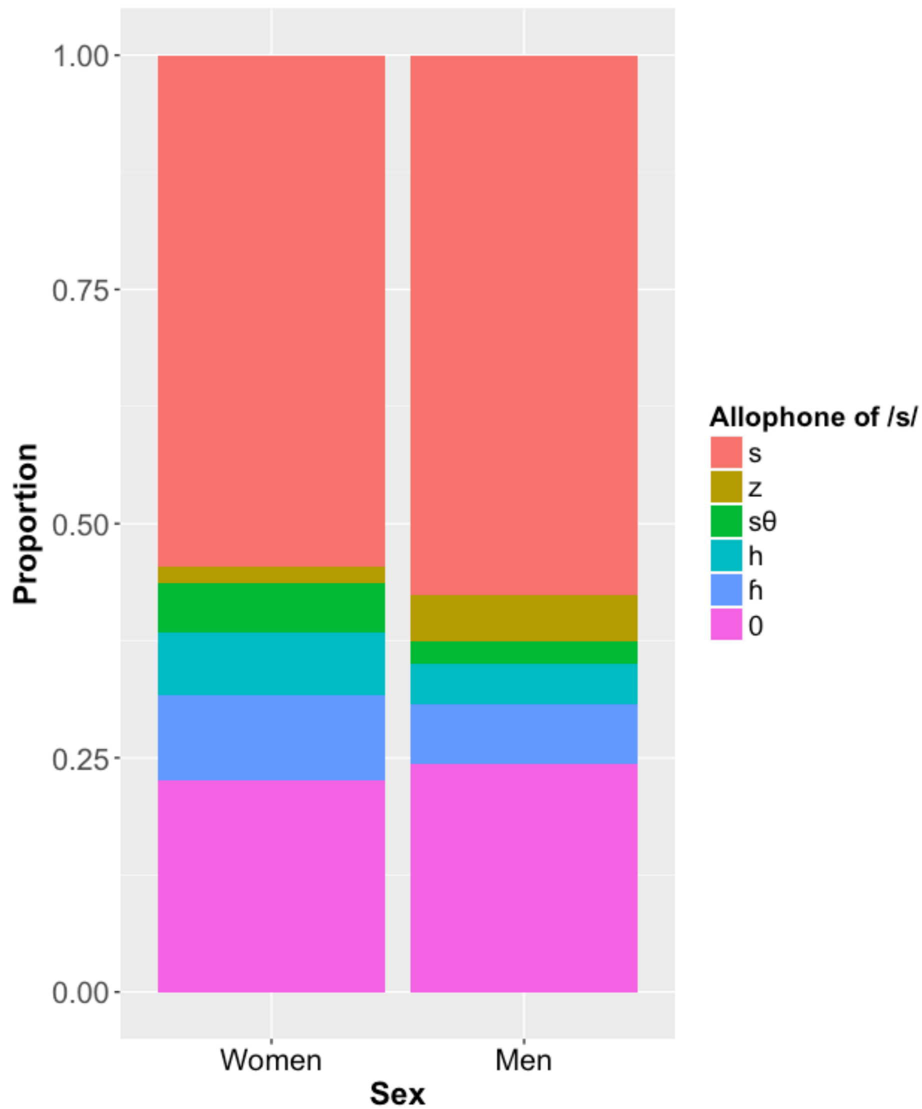

Figure 69. Proportion of allophones produced by sex.

Also evident in both Figure 68 and Figure 69 is the large disparity in the importance of constraint PRESERVE(segmental), which favors [ɦ] over full elision. While the less standard groups in this study have consistently shown a higher ranking of this constraint in previous sections, no difference has been as great as the one we see between men and women. This is because, as we see in Figure 69, women not only produce more [ɦ] than men, but they actually delete /s/ less often.

While group differences for other social factors have been more or less clear-cut, gender differences reveal that men and women have their own preferred variants of lenited /s/. It is true that women lenite more overall, but the difference in production of [s] is, in fact, quite small. The most striking difference is the prevalence of the approximant [s^{θ}] and debuccalized [h] and [ɦ] in the speech of women, while men produce more outputs containing [z] and [Ø]. The implications of these social differences, some of the most interesting thus far, are explored in Section 7.3.

In an effort to illustrate what these important differences mean for predicted probabilities in the grammars of men and women, let's turn our attention to Table 52 and Table 53, which compare constraint weights and resulting predicted probabilities for the word *casa* /kasa/ ('house') by gender. This syllable-initial, atonic prosodic position is particularly interesting to examine because the landscape of possible outcomes for each gender is considerably different.

Table 52

Tableau for /kasa/ ('house'); women

/kasa/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (strid, voice)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)	Harmony	Predicted	Observed
weight	3.20	1.94	2.22	3.29	0.00	5.13	0.81	4.13	0.00	5.28			
a. ˈka.sa	*	*	*	*	*						10.65	0.71	26
b. ˈka.za	*		*	*		*			*		13.84	0.03	0
c. ˈka.s^{θ}a	*	*	*			*	*				13.30	0.11	6
d. ˈka.ha	*	*				*	*	*			15.21	0.02	0
e. ˈka.ɦa	*					*	*	*	*		13.27	0.12	2
f. ˈka.∅a						*	*	*	*	*	15.35	0.01	0

Table 53

Tableau for /kasa/ ('house'); men

/kasa/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (strid, voice)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)	Harmony	Predicted	Observed
weight	4.26	1.20	1.72	1.03	0.73	3.63	0.58	2.67	0.00	5.53			
a. ˈka.sa	*	*	*	*	*						8.94	0.69	8
b. ˈka.za	*		*	*		*			*		10.64	0.13	1
c. ˈka.s^{θ}a	*	*	*			*	*				11.39	0.06	5
d. ˈka.ha	*	*				*	*	*			12.34	0.02	0
e. ˈka.ɦa	*					*	*	*	*		11.14	0.08	2
f. ˈka.∅a						*	*	*	*	*	12.41	0.02	0

Even though women, on average, lenite /s/ more readily than men, the relative importance of PRESERVE(strident, voice) in Table 52 as compared to Table 53 reveals that output [ˈka.sa] is slightly more probable in female speech than in male speech. This, combined with the relative importance of PRESERVE(coronal) and relative unimportance of PRESERVE(strident), renders [ˈka.s^{θ}a] 8% more likely than [ˈka.za] in the female grammar. In the male grammar, contrastingly, the relative importance of both PRESERVE(strident, voice) and PRESERVE(coronal) as compared to PRESERVE(strident) is diminished, and [ˈka.za] becomes the second-most likely output with a predicted probability of 13%. For both groups, production of a segment with some degree of oral constriction is by far the most likely scenario; however, this is less true for women, for whom [ˈka.ɦa] is also a fairly harmonic candidate.

6.5.5 Education

The final social factor we will examine in isolation is level of education. Education, while well-established as an important conditioning factor for sociophonetic variation generally and for Spanish /s/ weakening specifically, is somewhat problematic in this study because participants were not balanced (see Chapter 5 for a justification of this decision). Therefore, before examining education-graded differences in lenition patterns, it is important to look at the demographic composition of the two educational groups. Figure 70 shows the proportion of each educational group that pertains to each regional group, origin, age group, and gender.

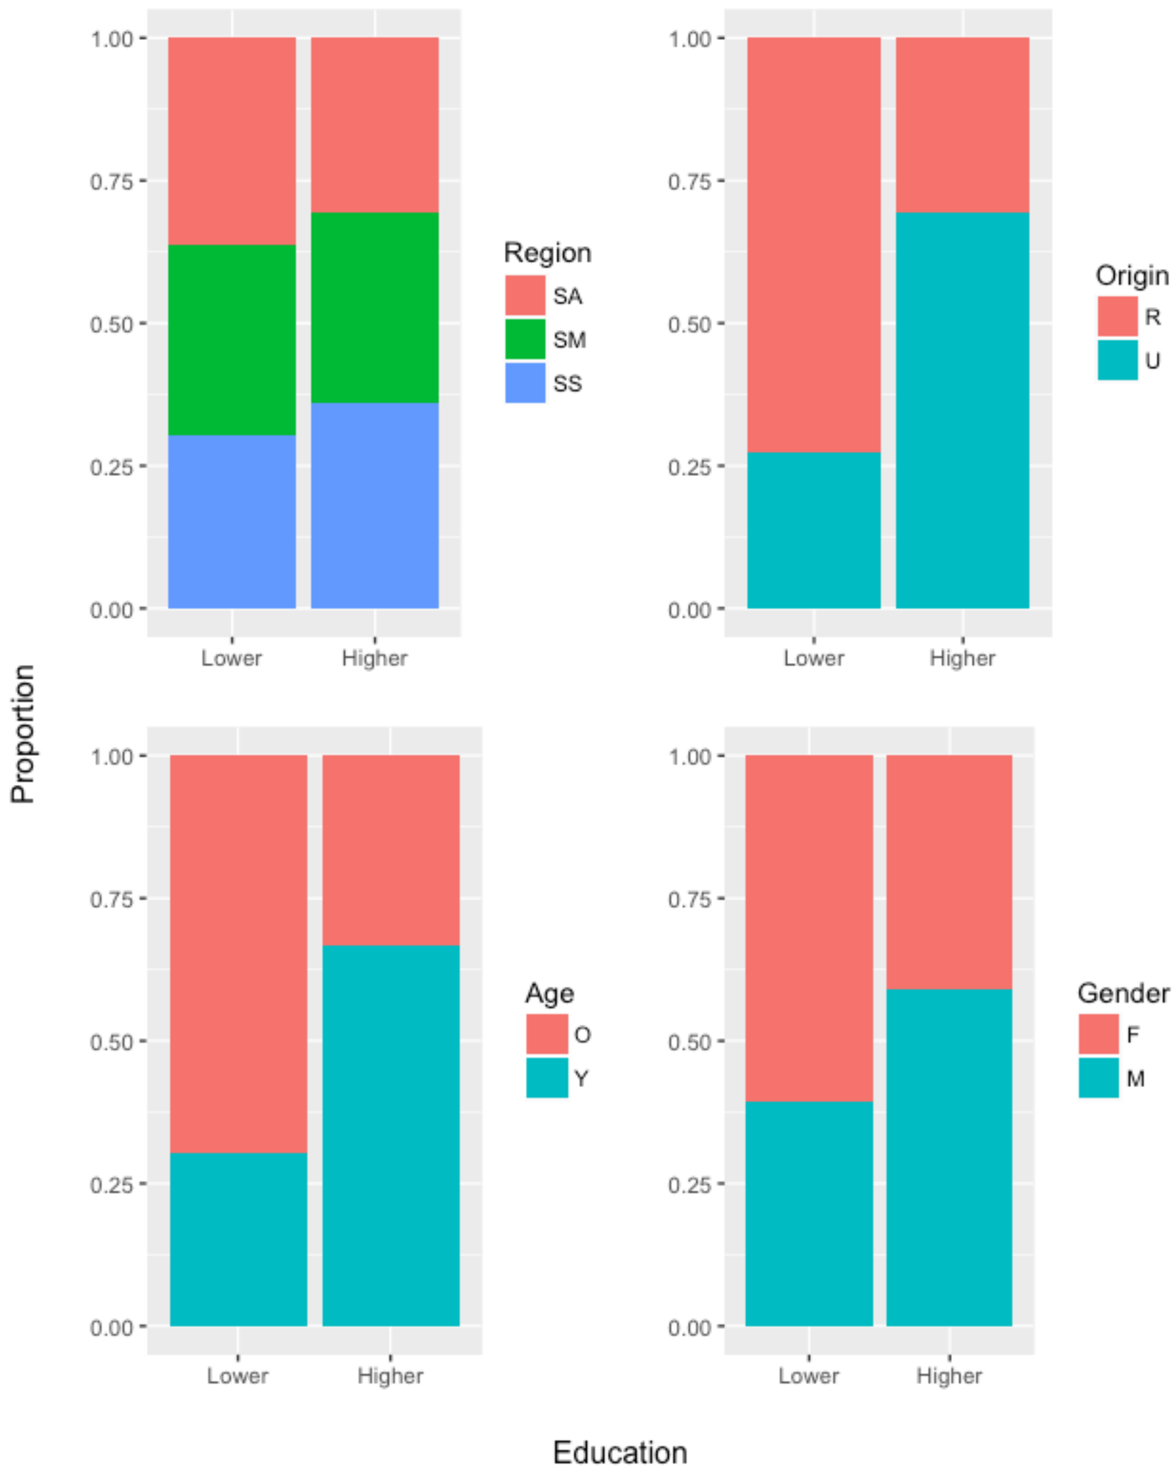

Figure 70. Demographic composition of the lower and higher educational groups.

Here we see that speakers from the lower educational group are disproportionately rural, older, and female. This isn't surprising; in fact, it is possible that a lower level of education is what unites these three social groups, who all pattern similarly with respect to /s/ lenition. While it is impossible to say whether the relationship between level of education and these three factors is correlation, causation, or simply random given the small sample size, this relationship is worth addressing and will be explored in the next chapter.

Given the observed relationship between level of education and factors of origin, age, and gender—in addition to the literature on the role of education in determining rates and patterns of /s/ lenition—it is unsurprising to see that the lower educational group weakens /s/ at much higher rates than the group with more education. This is evident when we look at general trends in educational differences in Table 54, which compares differences between Faithfulness constraint weights and Markedness constraint weights for the two educational groups as compared to the base grammar.

Table 54

Relative importance of Faithfulness by education

Group	Average Markedness constraint weight	Average Faithfulness constraint weight	Relative importance of Faithfulness	Scaling factor
El Salvador (all)	1.29	2.16	0.87	N/A
Lower	1.17	2.04	0.86	-0.01
Higher	1.07	2.08	1.01	+0.14

Table 54 confirms that Faithfulness is more important in the grammar of Salvadorans with more education; this observation is again confirmed in the raw data, visualized in Figure 71.

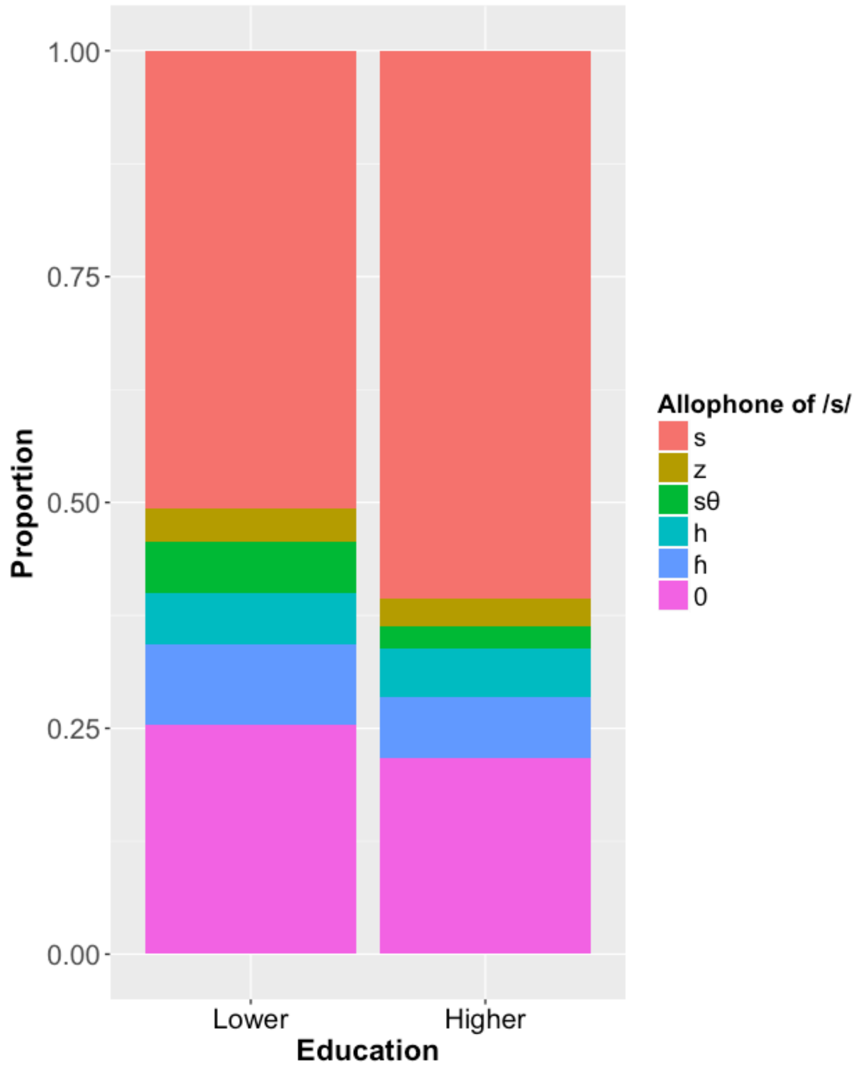

Figure 71. Proportion of allophones produced by level of education.

Participants with less education not only lenite /s/ at higher rates, but also employ more allophones outside the traditional tripartite system (i.e., [s^{θ}], [z], and [ɦ]). Furthermore, and unsurprisingly, speakers from the lower educational group lenite at much higher rates in more salient prosodic positions. In fact, as the reader can see in Table 55 and Table 56, speakers with less education are three times as likely to delete and two-thirds more likely to debuccalize /s/ in syllable-initial atonic position, exemplified by predicted probabilities for input *esa* /esa/ ('that'; feminine).

Table 55

Tableau for /esa/ ('that'; feminine); some primary education or less

/esa/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (strid, voice)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)	Harmony	Predicted	Observed
weight	3.82	2.18	2.32	3.14	0.00	5.01	0.00	3.21	0.00	5.46			
a. ˈe.sa	*	*	*	*	*						11.46	0.68	25
b. ˈe.za	*		*	*		*			*		14.29	0.01	5
c. ˈe.sᶿa	*	*	*			*	*				13.33	0.10	5
d. ˈe.ha	*	*				*	*	*			14.22	0.04	1
e. ˈe.ɦa	*					*	*	*	*		12.04	0.13	6
f. ˈe.∅a						*	*	*	*	*	13.68	0.03	0

Table 56

Tableau for /esa/ ('that'; feminine); some secondary education or more

/esa/	LAZY (any_ gesture)	LAZY (glot_ abduct)	LAZY (oral_ gesture)	LAZY (strid_ fric)	LAZY (vcl_ strid_fric)	PRES (strid, voice)	PRES (strid)	PRES (cor)	PRES (voice)	PRES (seg)	Harmony	Predicted	Observed
weight	3.93	1.62	1.88	1.40	0.21	4.14	0.00	3.04	0.00	5.89			
a. ˈe.sa	*	*	*	*	*						9.04	0.75	28
b. ˈe.za	*		*	*		*			*		11.35	0.07	2
c. ˈe.sᶿa	*	*	*			*	*				11.57	0.06	0
d. ˈe.ha	*	*				*	*	*			12.73	0.02	0
e. ˈe.ɦa	*					*	*	*	*		11.11	0.09	3
f. ˈe.∅a						*	*	*	*	*	13.07	0.01	0

As the reader can see, the lower relative importance of Faithfulness in Table 55 as compared to Table 56 allows for richer variation for the group of speakers with less education. In a prosodic position in which /s/ weakening is considered rare or nonexistent in most dialects of Spanish, there is a one-in-three chance that speakers will produce a non-strident allophone, and a one-in-five chance that that allophone will require no oral constriction at all. Contrastingly, for speakers with more education, the probability that the feature [+strident] will be lost in the output in this prosodic position is less than one-in-five, and the probability that the feature [+coronal] will be lost is approximately one-in-ten. It should also be noted that the predicted probability of deletion in this position by the less educated group is the highest of any group presented in this section at 3%.

6.5.6 Interactions between social groups

Thus far, Section 6.5 has compared the relative importance of Faithfulness in the grammar according to five social factors: region, origin, age, gender, and level of education. To conclude this exploration, this final subsection examines significant interactions between social groups that have been discussed in this chapter. In the interest of space, I opted to model these interactions using multinomial logistic regression in the 'nnet' package in R (Venables & Ripley, 2002) instead of maximum entropy; however, as described in Chapter 5, both methods rely on gradient ascent/descent and the underlying calculations are comparable.

Because the relationships between level of education and the other four social factors were explored in the previous section, here I examined only interactions between region, origin, age, and gender. Before constructing the multinomial regression, I explored the raw data visually in order to identify the most interesting interactions to model. Figure 71 below shows two-way relationships between the four social factors of interest and proportion of allophone produced.

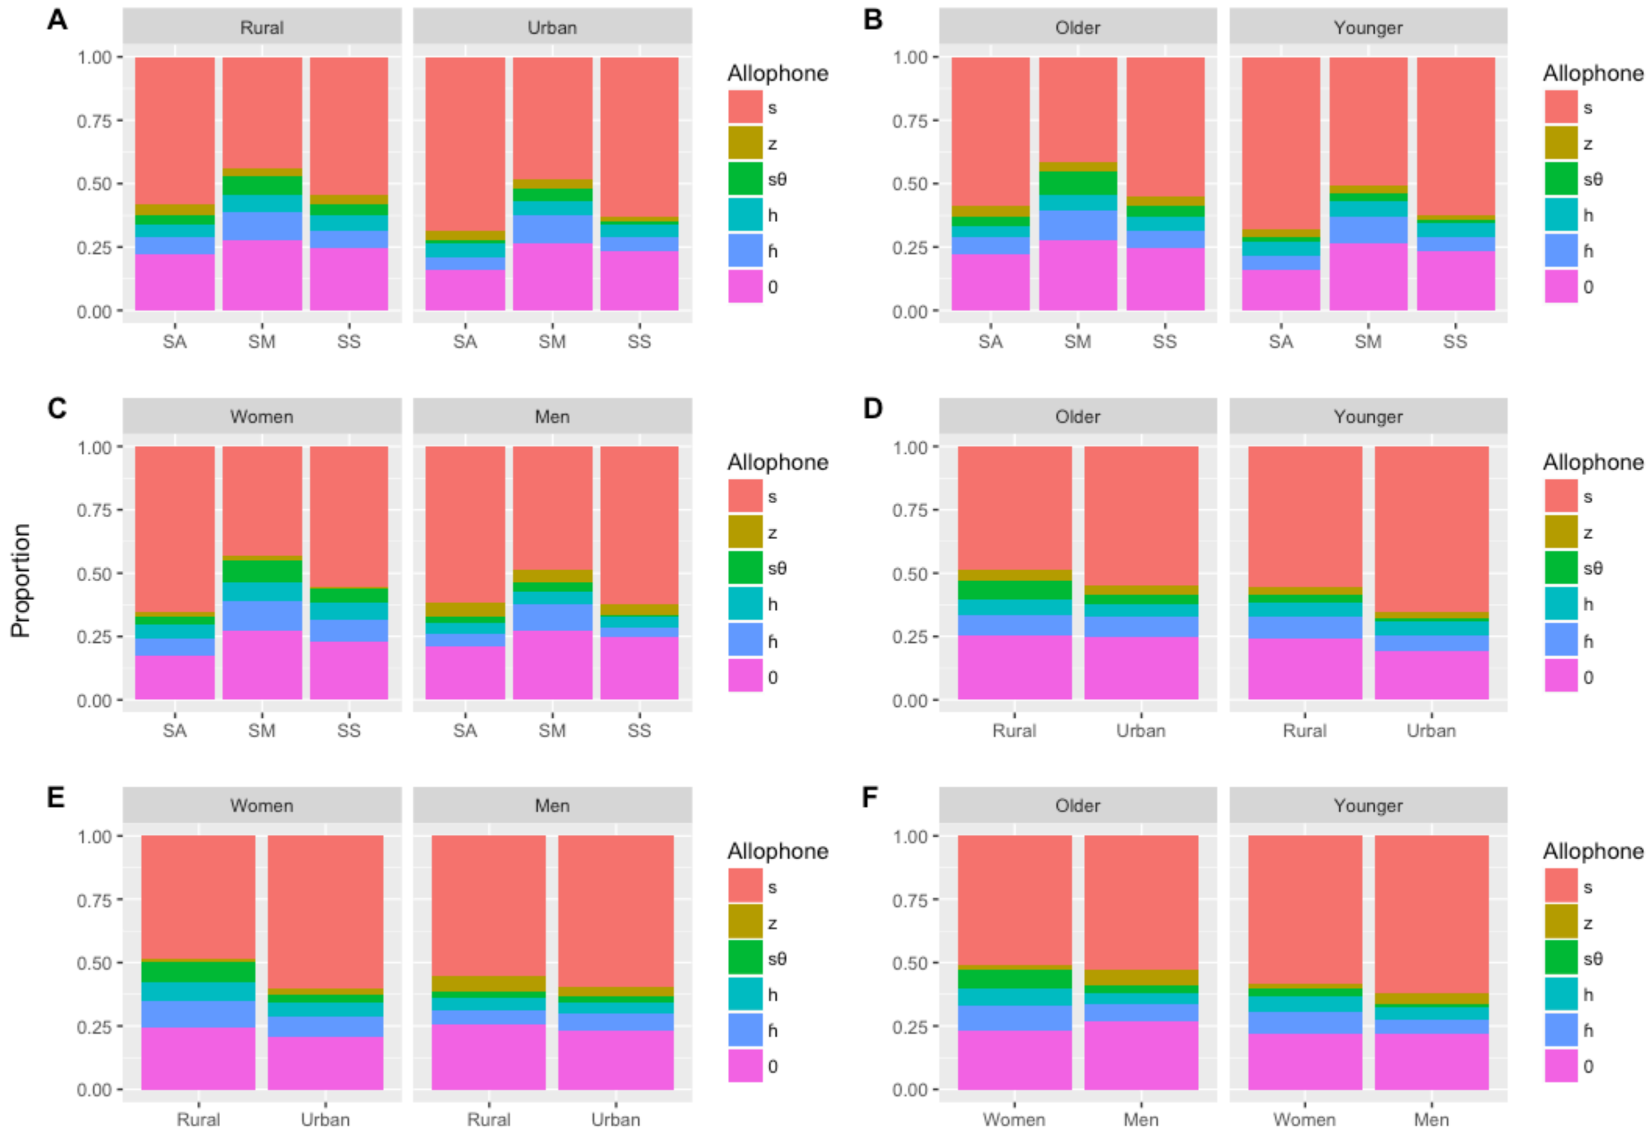

Figure 72. Relationships between region and origin (A), region and age (B), region and gender (C), origin and age (D), origin and gender (E), and gender and age (F) with respect to proportion of allophone produced.

Figure 71 reveals a number of interesting interactions in the data. In Figure 71a, we see that, while origin has an effect on allophonic production for speakers from Santa Ana and San Salvador, it appears to matter much less for speakers from San Miguel, who lenite at high rates regardless of whether they're from a rural or urban community. Interestingly, in Figure 71b, we see that the interaction of age and region is almost identical to that seen in Figure 71a: age determines allophonic variation for speakers from the more standardized regions, but not in San Miguel (with the exception of [s^{θ}], which is produced at much higher rates by older San Migueleños). In Figure 71c, the most striking differences between the male and female groups is the use of [s^{θ}], which is particularly prevalent in the speech of women from San Miguel and San Salvador but seems to be exchanged for [z] in male speech. The interaction between region and gender is particularly interesting because, while Salvadoran women as a whole lenite more than their male counterparts, the women from Santa Ana appear to be the most conservative of any group compared in Figure 71c. In Figure 71d, we see that groups align as we might expect: older rural speakers lenite at the highest rates and younger urban speakers lenite the least, with younger rural speakers and older urban speakers producing almost identical allophonic variation. Finally, in Figure 71e we see that origin matters much more for women than for men, while in Figure 71f we see that age interacts with both genders to a similar degree.

Based on these results, I chose to model the interactions between region and age, region and gender, and gender and origin. Because interpreting interaction effects of categorical variables in a multinomial regression depends in part on the reference levels chosen for the variables in question, I selected levels with the least amount of observed within-category variation in Figures 71a-e: 'San Miguel' for region, 'urban' for origin, 'younger' for age, and 'men' for gender. The results of the model are presented in Table 57.

Table 57

Results of multinomial logistic regression for interaction effects

Allophone	Coefficient	SE	*p* value
Intercept			
[z]	-2.73	0.15	<0.001***
[s^{θ}]	-3.46	0.17	<0.001***
[h]	-2.49	0.13	<0.001***
[ɦ]	-1.68	0.10	<0.001***
[Ø]	-0.79	0.07	<0.001***
Region x Age (San Salvador, Older)			
[z]	0.58	0.25	<0.05*
[s^{θ}]	-0.30	0.24	0.22
[h]	-0.08	0.18	0.66
[ɦ]	-0.05	0.16	0.73
[Ø]	-0.05	0.10	0.62
Region x Age (Santa Ana, Older)			
[z]	0.27	0.23	0.23
[s^{θ}]	-0.67	0.24	<0.01**
[h]	-0.32	0.18	0.08
[ɦ]	-0.01	0.16	0.95
[Ø]	0.25	0.10	<0.05*
Region x Gender (San Salvador, Women)			
[z]	-0.58	0.28	<0.05*
[s^{θ}]	1.13	0.29	<0.001***
[h]	0.22	0.19	0.24
[ɦ]	0.58	0.16	<0.001***
[Ø]	0.08	0.10	0.43
Region x Gender (Santa Ana, Women)			
[z]	-0.44	0.25	0.08
[s^{θ}]	-0.74	0.23	<0.01**
[h]	-0.41	0.19	<0.05*
[ɦ]	-0.03	0.16	0.87
[Ø]	-0.38	0.10	<0.001***
Gender x Origin (Women, Rural)			
[z]	-1.03	0.23	<0.001***
[s^{θ}]	1.08	0.20	<0.001***
[h]	0.35	0.15	<0.05*
[ɦ]	0.66	0.13	<0.001***
[Ø]	0.21	0.08	<0.05*

In Table 57, the coefficient of each interaction effect represents the change in the effect of group A when group B changes. For example, the interaction coefficients for gender by origin compare the effects of switching from urban to rural for women with the same change for men. A positive coefficient signifies that the effect for the group listed is larger than for the reference group; a negative coefficient signifies the opposite.

A number of significant interaction effects were obtained in the multinomial logistic regression model. With respect to the effect of region by age, the relative log odds of producing [z] increase significantly more between younger and older speakers in San Salvador than in San Miguel ($p < 0.05$). In other words, age matters more for production of [z] in the former region than it does in the latter. Additionally, the relative log odds of producing [s^{θ}] increase significantly less between younger and older speakers in Santa Ana than between younger and older speakers in San Miguel ($p < 0.01$), while the relative log odds of producing [∅] increase significantly more between younger and older speakers in Santa Ana than in San Miguel ($p < 0.05$). The interaction between region and age can be visualized in Figure 73, which shows predicted probabilities calculated by the regression model for allophones [z], [s^{θ}], and [∅].

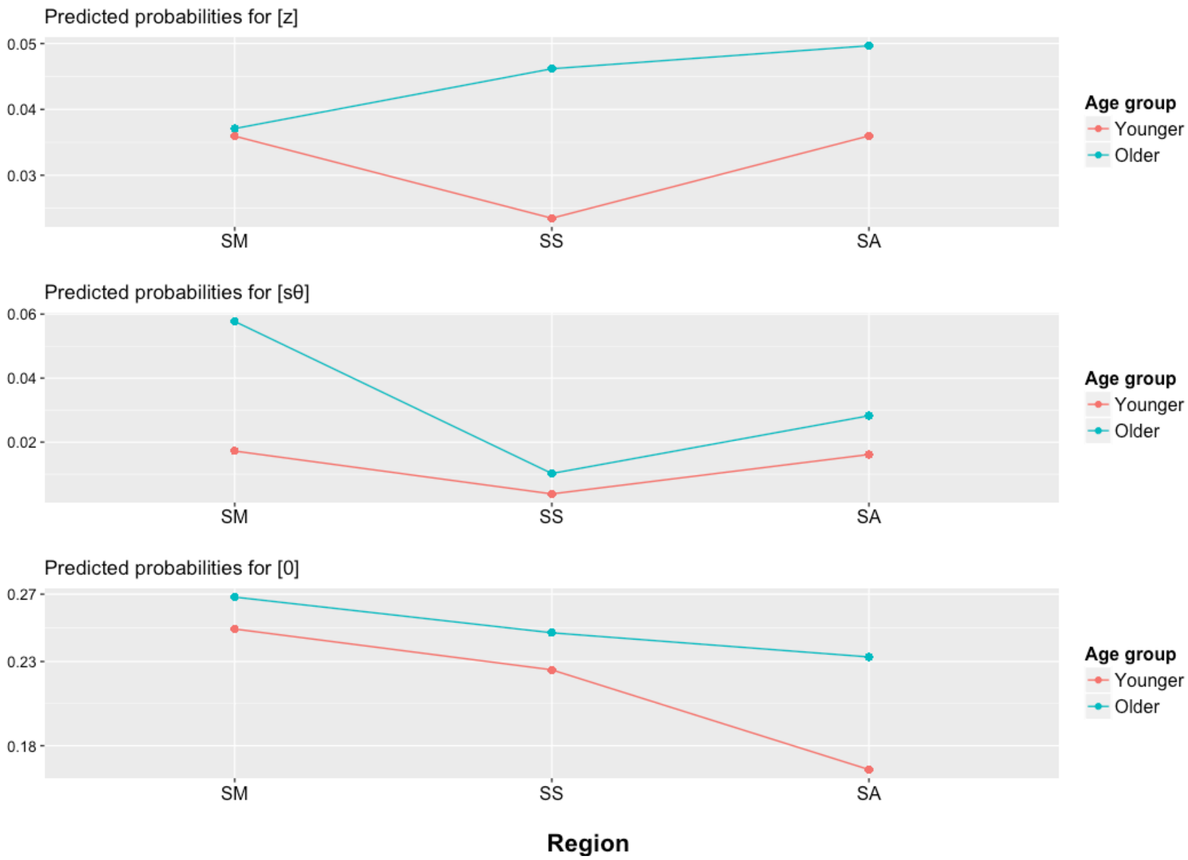

Figure 73. Predicted probabilities for allophones [z], [$s^θ$], and [∅] by region and age group.

With respect to the effect of region by gender, significant effects in the model confirm the initial observation that women from Santa Ana are surprisingly standard. In Figure 74, we see that the relative log odds of producing [$s^θ$] increases for women in San Salvador yet decreases for women in Santa Ana. The same is true of [∅], albeit to a much lesser degree. With respect to [h], the effect for gender within San Miguel and Santa Ana is in the same direction, but the relative log odds of producing this allophone increase significantly more in San Miguel. Additionally, the relative log odds of producing [z] decrease significantly more for women in San Salvador than in San Miguel, yet the relative log odds of producing [$s^θ$] increase significantly more for women in San Salvador than in San Miguel. Finally, with respect to [ɦ], we see in Figure 74 that the log odds of producing [ɦ] increase for women in San Salvador but decrease slightly for women in San Miguel.

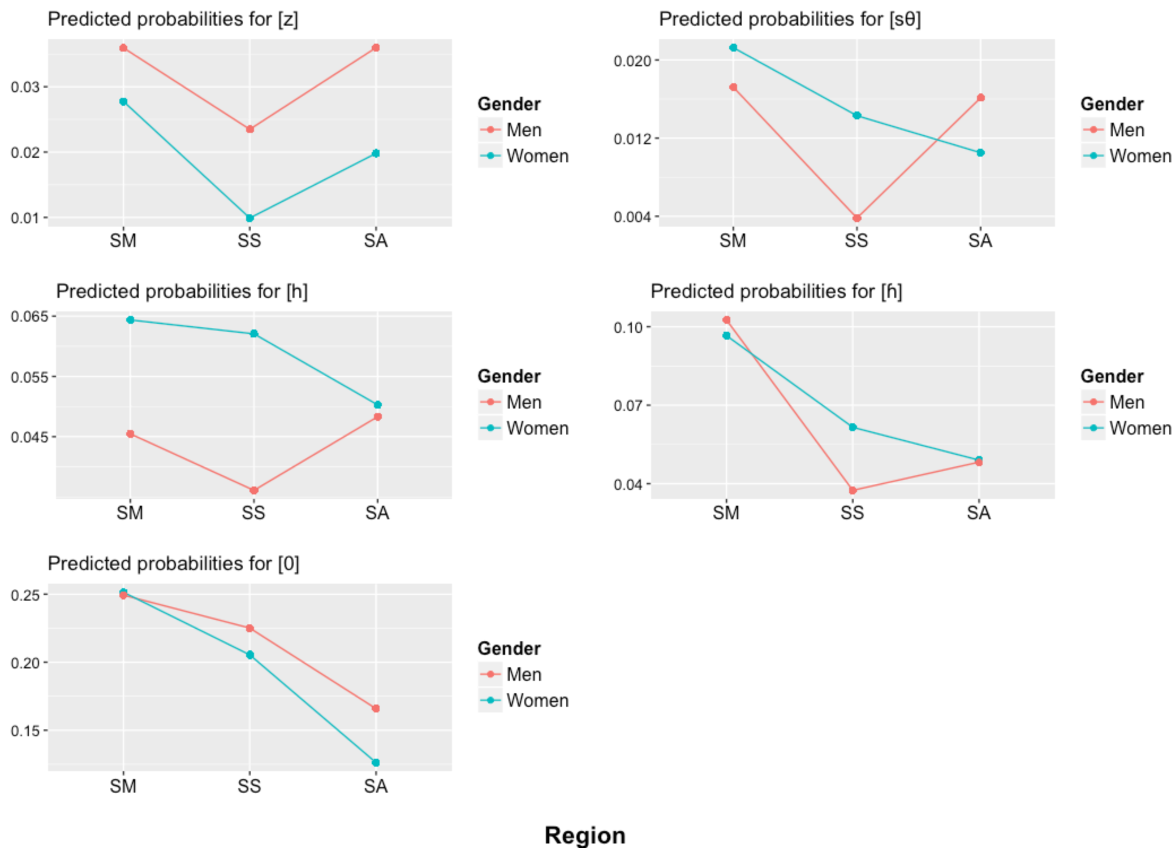

Figure 74. Predicted probabilities for all non-[s] allophones by region and gender.

Finally, significant effects were obtained for all five allophones for the interaction of gender by origin. With respect to [z], we see that rural men are more likely to produce [z] than urban men while rural women are less likely to produce [z] than urban women (although use of [z] among women generally is quite low). Contrastingly, urban men are more likely to produce [ɦ] than rural men, while urban women are less likely to produce [ɦ] than rural women. Effects for allophones [s^{θ}], [h], and [∅] are in the same direction for men and women, but the relative log odds of producing these allophones increase significantly more for rural women than for rural men.

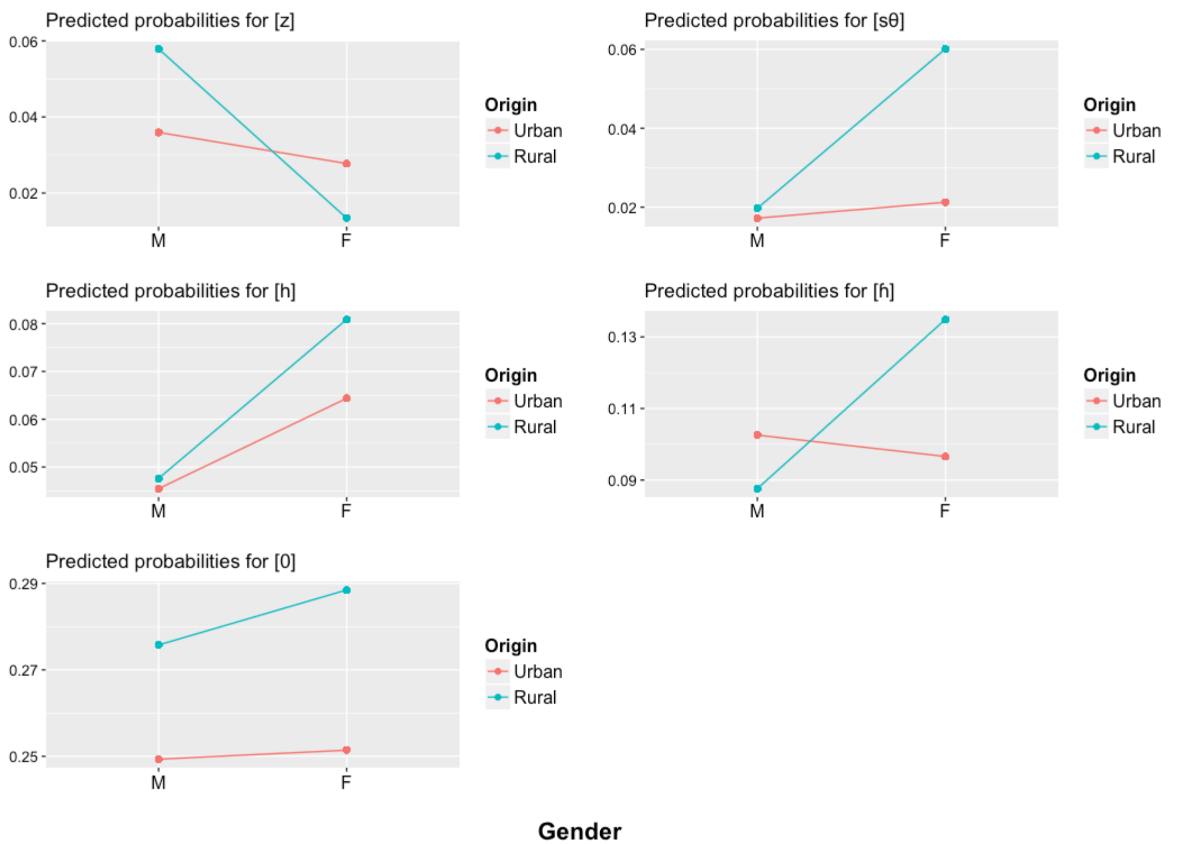

Figure 75. Predicted probabilities for all non-[s] allophones by gender and origin.

This subsection has modeled interactions between non-grammatical variables as part of a larger section on the role of social factors in modulating patterns of /s/ lenition. While the linguistically-based model presented in section 6.4 is able to capture overarching and fundamental tendencies of the grammar, there is some amount of variation that cannot be reconciled without considering the social characteristics of speakers. This section has identified San Migueleños, those from rural communities, older Salvadorans, and women as the groups most likely to lenite /s/ generally, to use particularly non-standard allophones such as [s^{θ}], and to weaken the voiceless sibilant in more salient prosodic positions. Possible interpretations of these results, which must rely on both the sociolinguistics canon and sociopolitical/historical factors unique to El Salvador, will be presented in the following chapter.

6.6 Conclusion

This chapter has achieved three principle goals. First, and most importantly, it has provided a comprehensive analysis of Salvadoran /s/ weakening via the interaction of effort-based Markedness and perceptually-based Faithfulness constraints. This analysis has been couched within an Optimality-Theoretic framework but has proposed weighting constraints via maximum entropy in lieu of strict ranking, which is unable to account for the quasi-free variation observed in the Salvadoran data. Results show that an effort-based approach in the style of Kirchner (1998, 2004) does a good job of accounting for variation in the data, although potential shortcomings of this account will be discussed in the next chapter.

Second, this chapter has presented various sub-models of the maxent grammar in which the relative importance of Faithfulness is scaled up or down according to social groups pertaining to different regions, places of origin, age groups, genders, and levels of education. The use of a 'scaling factor' has allowed us to see where the grammar and its corresponding predictions make adjustments for different categories of speakers and has demonstrated that all Salvadoran speakers adhere to a common base grammar despite assigning slightly different weights to the constraints within it.

Finally, this chapter has explored the effects of both lexical frequency and social factor interactions using statistical methods separate from the maxent grammar. Lexical frequency did not prove to play a significant role in determining patterns of /s/ lenition in the Salvadoran data and was therefore excluded from the final model. Contrastingly, social factor interactions were numerous, but were modeled using multinomial logistic regression instead of maxent in the interest of space and feasibility. Chapter 7 discusses the significance of these statistical models as well as those learned by the Maxent Grammar Tool in relation to the theories presented in chapters 2 and 3 with the goal of addressing the research questions and hypotheses at the core of this dissertation.

Chapter 7. Discussion & conclusion

7.1 Introduction

Chapter 6 presented the maxent base grammar and explored effects on the relative importance of Faithfulness constraint weights when various social groups were modeled individually. In the present chapter, these results are analyzed in more depth and important connections are drawn between the Salvadoran data and the existing literature and theoretical frameworks presented in Chapters 2 and 3, respectively. Chapter 7 begins with a discussion of Markedness and Faithfulness in the grammar: Section 7.2 synthesizes important findings and digs deeper into the data to provide interpretations within the context of phonological and phonetic theory. Then, Section 7.3 does the same for social variation, comparing results to previous work on Spanish /s/ in other dialects and introducing a wealth of theory and demographic statistics to contextualize unexpected findings.

The chapter then switches gears, transitioning from explicit interpretations of the results to their broader implications. Section 7.4 compares the 2018 Salvadoran data with lenition rates reported in Lipski (1984) in an effort to both forge some kind of longitudinal comparison and to highlight the need for longitudinal work with comparable population samples. Next, Section 7.5 addresses issues posed in Chapter 2 regarding pathways of diffusion in dialects with both onset and coda weakening, and Section 7.6 discusses the lack of lexical frequency effects in the data and offers possible explanations. Finally, Section 7.7 offers a comprehensive conclusion to this dissertation, and Section 7.8 proposes future directions of study.

7.2 The maxent grammar

In Chapter 6 we saw that the interactions of constraints that reflect the need to minimize articulatory effort cost while preserving salient perceptual distinctions are able to account for important patterns in the Salvadoran /s/ lenition data. This section synthesizes findings with respect to Markedness and Faithfulness constraints and explores possible theoretical and methodological explanations for findings that do not align with original hypotheses.

7.2.1 Markedness constraint weights and implications for the grammar

Chapter 6, Section 6.4.1 presented findings for Markedness constraint weights according to constraint and context. With respect to intervocalic contexts, we saw that hypotheses were strongly confirmed: when flanking segments are more open, the biomechanical energy needed to move the tongue to the articulatory target increases and the impetus to lenite /s/ increases in turn. As expected, weights for LAZY(vcl_strid_fric) and LAZY(strid_fric) were identical for two of the intervocalic contexts and approximately the same for the other two, while weights for LAZY(oral_gesture) were distinctly lower. Furthermore, as predicted, weights for LAZY(glottal_abduction) did not differ significantly based on adjacent vowel height as the openness of flanking segments should have no effect on the effort cost of separating the vocal folds. While accounts of /s/ weakening in other dialects of Spanish have tended to conceive of preceding and following segments as separate entities instead of taking them together as a single environment, my results nevertheless echo previous findings that a preceding or following non-high vowel favors lenition more than an adjacent high vowel (e.g., Esther Brown, 2005; Earl Brown & Esther Brown, 2012; Esther Brown & Torres Cacoullos, 2002, 2003).

While I did not offer any hypotheses regarding lenition in $V_{[-high]}$ __ $V_{[+high]}$ and $V_{[+high]}$ __ $V_{[-high]}$ contexts beyond these contexts favoring weakening more than $V_{[+high]}$ __ $V_{[+high]}$ and less than $V_{[-high]}$ __ $V_{[-high]}$, the model revealed that weakening is favored slightly more when /s/ is followed by a vowel that is [-high]. Upon exploring the raw data, it appears that the effort cost of producing an oral gesture is more dependent on the adjacent vowel that pertains to the same syllable as /s/, which can be seen in Figure 76.

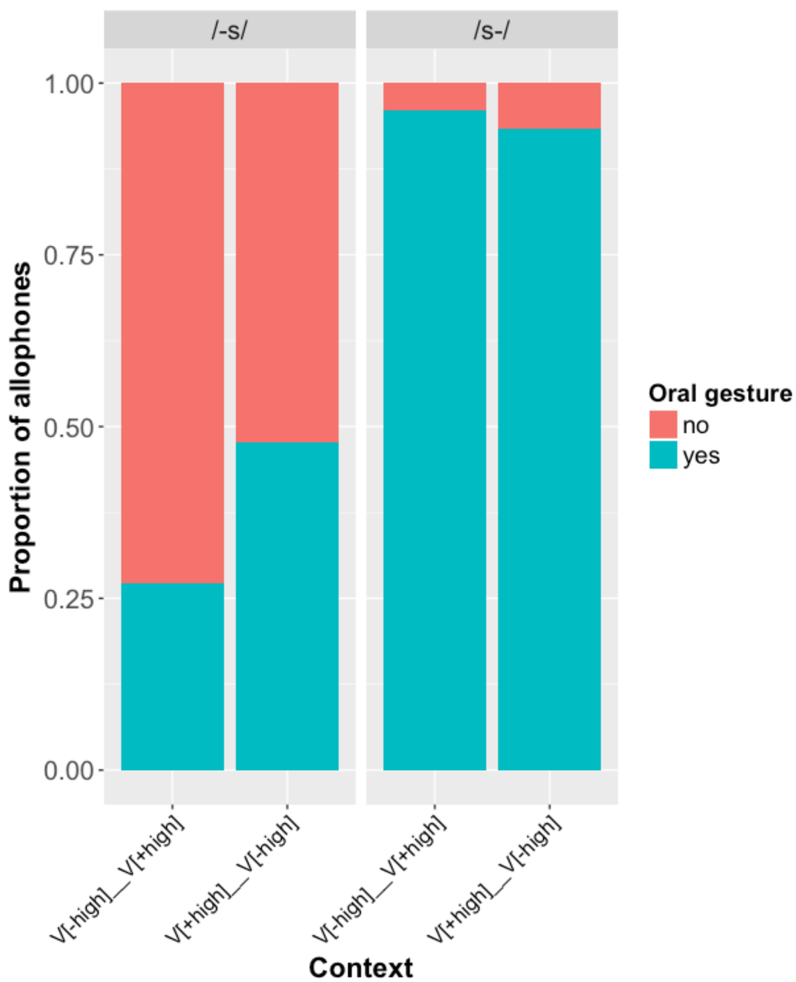

Figure 76. Proportion of allophones that required an oral gesture in $V_{[-high]}$ __ $V_{[+high]}$ and $V_{[+high]}$ __ $V_{[-high]}$ contexts by syllable position.

Here we see that coda /-s/ (*n* = 331) is more likely to debuccalize or delete when the non-high vowel precedes it and onset /s-/ (*n* = 1915) is more likely to debuccalize or delete when the non-high vowel follows. A binary logistic regression model confirmed that these effects are significant for both coda /-s/ ($p < 0.05$) and onset /s-/ ($p < 0.05$). While it is logical that the vowel pertaining to the same syllable as /s/ would play a larger role in determining patterns of weakening, this result contradicts Earl Brown and Esther Brown (2012, pp. 105–106), who argue that /s-/ weakening is more sensitive to preceding non-high vowels while /-s/ weakening is more sensitive to following non-high vowels. It is possible that this discrepancy is due to dialectal differences, but further research into this line of inquiry is needed.

With respect to preconsonantal contexts, LAZY(vcl_strid_fric) is logically weighted highest in contexts that are perceived to be the most difficult for [s]: when the preceding vowel is more open and the following consonant disagrees with [s] in both coronality and voicing, as in the word *rasgo* /rasgo/ (‘trait’), and lowest in contexts that have the least effort cost: between two consonants as in *instante* /instante/ (‘instant’) and between a closed vowel and a consonant that agrees with [s] in both coronality and voicing as in *listo* /listo/ (‘ready/smart’). However, for the five other preconsonantal contexts, it appears that coronality of the following consonant is the most important factor, voicing plays a secondary role, and the height of the preceding vowel doesn’t matter. Figure 77 illustrates this perceived effort scale based on constraint weights learned by maxent.

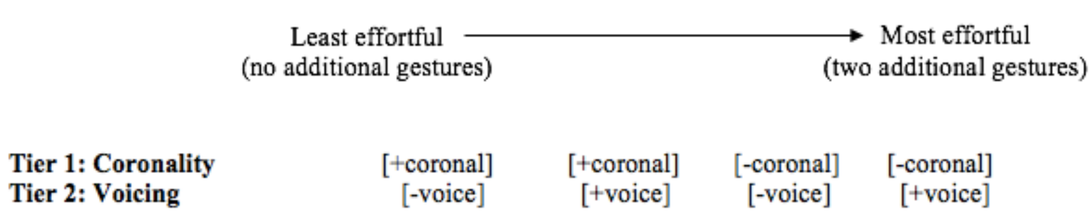

Figure 77. Perceived effort scale in V_C contexts based on coronality and voicing of the following consonant.

The role that the openness of the preceding segment plays in preconsonantal contexts is difficult to ascertain. At the poles of the "effortfulness" spectrum it appears that more closed preceding segments favor retention while more open preceding segments favor lenition; however, this difference is inconsistent for the intermediate contexts. It is possible that, if a context is judged to be either very easy (no additional gestures) or very difficult (two additional gestures), the role of the preceding vowel is simply to distinguish between easy (/esta/) and easiest (/listo/) and between hard (/mismo/) and hardest (/rasgo/).

As discussed in Chapter 6, weights for LAZY(strid_fric) and LAZY(glottal_abduction) in preconsonantal position reflect original hypotheses: the former patterns with LAZY(vcl_strid_fric) except when the following consonant is voiced and the effort cost (and therefore constraint weight) is reduced, and the latter is exclusively determined by the voicing of the following consonant, with mismatched glottal gestures incurring higher penalties. LAZY(oral_gesture), contrastingly, behaves predictably except in two contexts in which its weights are strikingly low compared to LAZY(vcl_strid_fric) and LAZY(strid_fric): $V_{[-high]}$__$C_{[-coronal,\ -voice]}$ and $V_{[+high]}$__$C_{[+coronal,\ +voice]}$. The discrepancies in constraint weights suggest that the effort cost for producing some oral gesture in these contexts is much less than for producing the precise, controlled gesture required for strident fricatives. For the first context this is entirely logical: in a word such as *aspecto* /aspekto/ ('aspect'), the distance the tongue must travel from low [a] to approximant [s^{θ}] is much less than if the active articulator must closely approach the alveolar ridge and remain there for enough time to produce an audible [s] before beginning the bilabial [p] gesture. For $V_{[+high]}$__$C_{[+coronal,\ +voice]}$, however, this effect is likely due to the small number of occurrences of this context in the data (n = 14). Like any machine learning tool, maxent's ability to capture patterns in the data depends, in part, on how much data it sees.

Because Markedness constraint weights for C_V contexts conformed to hypotheses and are well accounted for within an effort-based theory, the final context to discuss is pause-adjacent. In this context, as with the other three, the height of the preceding or following vowel has a systematic effect on constraint weights for LAZY(vcl_strid_fric), LAZY(strid_fric), and LAZY(oral_gesture), further confirming the role of effort cost in determining patterns of /s/ lenition. While the location of the pause (preceding or following /s/) appears to matter less than adjacent vowel height, it is worth examining the $V_{[+high]}$__## context as it appears to be markedly easier than the other three. Indeed, if we consider this context articulatorily, the tongue body in a word such as *país* /pais/ ('country') is already high from the preceding /i/ and primed to produce an [s]. Whereas in V_V or V_C contexts the speaker might be deterred from producing the voiceless sibilant due to the openness of the following vowel or dissimilarity of the following consonant, pre-pausal context leaves only the one adjacent segment to consider. As I discuss in Chapter 2, accounts of /s/ weakening in other dialects of Spanish have been varied in their reports of pre-pausal lenition patterns, with some reporting high rates of deletion and others reporting high rates of sibilant retention. Because these studies, to my knowledge, do not consider the height of the vowel preceding /s/ in these contexts, it is possible that the observed inconsistencies are in part due to contextual asymmetries (i.e., more cases of $V_{[+high]}$__## than $V_{[-high]}$__## or vice versa) in the data.[36]

7.2.2 Faithfulness constraint weights and implications for the grammar

Chapter 6, Section 6.4.2 presented findings for Faithfulness constraint weights according to phonological features of /s/ that should be preserved in particular prosodic positions in order to

[36] Another interesting consideration is that perhaps these asymmetries interact with morphology, as the plural morpheme /-s/ disproportionately occurs in $V_{[-high]}$ __ contexts.

maintain important perceptual distinctions. The constraint PRESERVE(strident, voice), which penalizes any output form that does not preserve input specifications for both [+strident] and [-voice], serves as the base constraint in the grammar because it is violated by any non-[s] candidate. Weights for this constraint reflected the initial hypothesis that the stronger the prosodic position in which /s/ occurs, the more resistant it will be to weakening. My hypothesis that phrase-final /s/, which occurs in word-final prosodic position (weak) but at the edge of a phrase boundary (salient), would prove to be more resistant to weakening than other coda /s/s but less resistant than onset /s/s was also confirmed.

Interestingly, however, results regarding the role of syllable stress were not as clear. That is, while tonicity appears to matter in phrase-initial, word-initial, word-final, and phrase-final positions, these same differences were not obtained in word-medial (i.e., syllable-initial and syllable-final) positions. These asymmetries suggest that, when /s/ occurs within a word instead of at its edges, syllable stress does not matter; this echoes Esther Brown and Torres Cacoullos's (2002, 2003) account of onset weakening in Chihuahuan Spanish in which the effect of tonicity was only significant word-initially. These findings might be accounted for by increased uncertainty on the part of the listener across phonological words. As Keating (2006, p. 180) explains, "from the perspective of the listener, initial segments are probably on average less determined by prior context and…therefore the acoustic signal must bear a greater load in the recovery of the message in these positions."

Within this framework it is possible that, in word-initial positions, syllable stress plays a more important role in preserving perceptual distinctions because the risk of confusability is higher. If this were the case, we would expect even within-category acoustic cues for [s] to be more modulated by tonicity in word-initial position than in syllable-initial position. Figure 78 compares word-initial and syllable-initial tokens of [s] by tonicity with respect to segment duration, which is a crucial component of prominence.

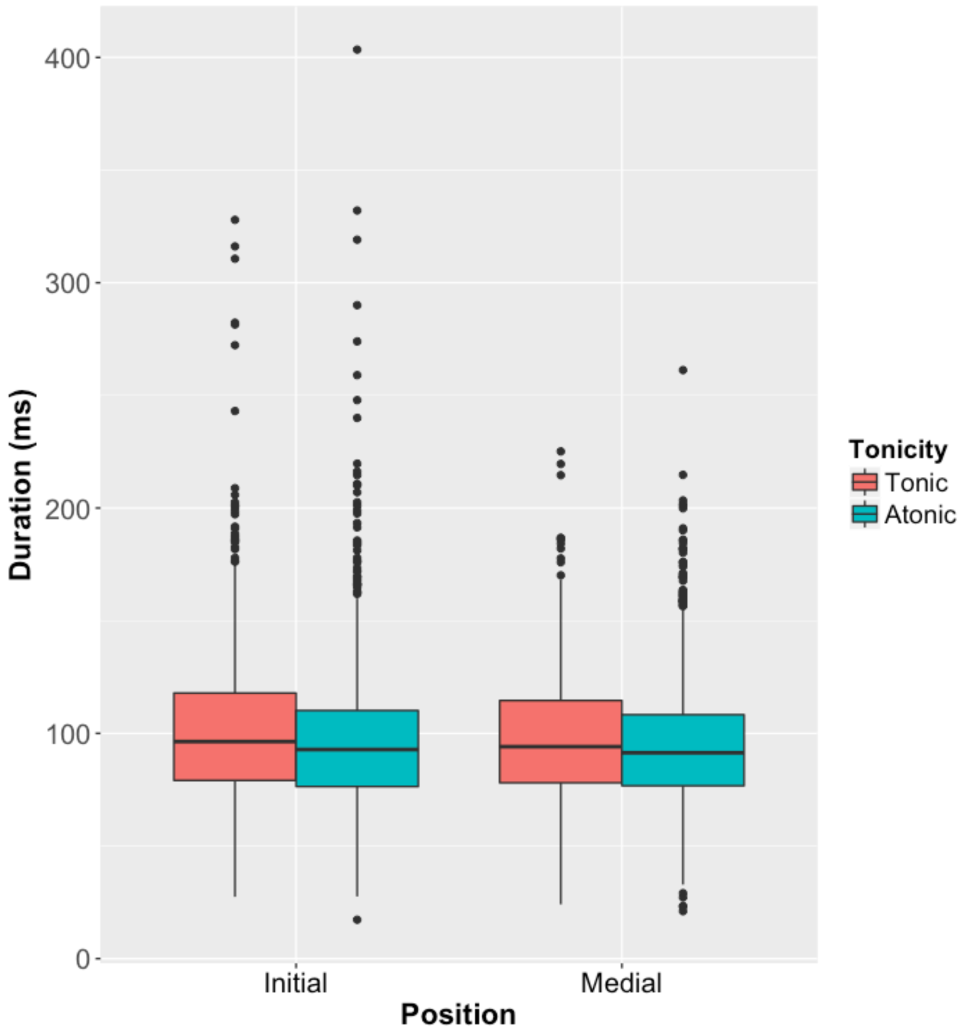

Figure 78. Comparison of duration measures for tokens of onset [s] occurring in word-initial and word-medial positions by tonicity.

In Figure 78, we see that duration measures for [s] are more affected by tonicity when the segment occurs in word-initial onset position than when it occurs in word-medial onset position. A linear regression model confirms that differences for duration measures in tonic and atonic word-initial positions are larger than differences in tonic and atonic word-medial positions, although both differences are statistically significant. Table 58 and Table 59 show the results of this model, with the first comparing duration measures between tokens of [s] in initial tonic and atonic positions and the second between medial tonic and atonic positions. In order to clearly present these comparisons, the level of the variable 'position' was releveled to obtain the results presented in Table 59.

Table 58

Linear regression with position and tonicity predicting duration of [s]

Position	Estimate	Standard error	*p* value
Intercept (Initial, tonic)	101.28	0.91	<0.001***
Initial, atonic	-5.22	1.21	<0.001***

Table 59

Linear regression with position and tonicity predicting duration of [s], releveled

Position	Estimate	Standard error	*p* value
Intercept (Medial, tonic)	97.23	0.84	<0.001***
Medial, atonic	-2.92	1.06	<0.01**

A comparison of the data presented in Table 58 and Table 59 shows that differences between durations of [s] in initial tonic and initial atonic positions are larger than differences in medial tonic and medial atonic positions, evident upon examining both the estimate values for the atonic positions as well as their respective *p* values. These results confirm that the distinction between stressed and unstressed syllables is more meaningful word-initially than word-medially: among tokens for which /s/ is not weakened allophonically, [s] undergoes more tonicity-based strengthening in the former position.

With respect to Faithfulness constraints that impose additional penalties for changes in input specifications for specific features, we observed a number of different patterns in Section 6.4.2. First, we saw that additional penalties for changing the input specification for [+strident] are highest in the three coda positions yet there are virtually no additional penalties in onset positions. This suggests that [z] and [s^{θ}] are equally bad as compared to [s] in onset positions, but [z] is better than [s^{θ}] in coda positions. Because the comparative strength of these two allophones is unclear (each violates two Faithfulness constraints) and because neither occurs frequently in coda positions, it is possible that these differences do not reflect prosodic strength relations but rather an accidental effect of Markedness. That is, the perceived preference for [z] in coda positions might be entirely due to calculations of effort cost. Indeed, a careful examination of the contexts in which coda [z] and [s^{θ}] occur in the raw data reveals that the former is favored before high vowels and voiced consonants while the latter is preferred before non-high vowels and voiceless consonants. This is logical with respect to effort cost as the higher tongue body required for [z] is easier to achieve before a high vowel than before a non-high vowel, and the glottal abduction required for [s^{θ}] is easier to achieve before a voiceless consonant. This comparison can be seen in Figure 79.

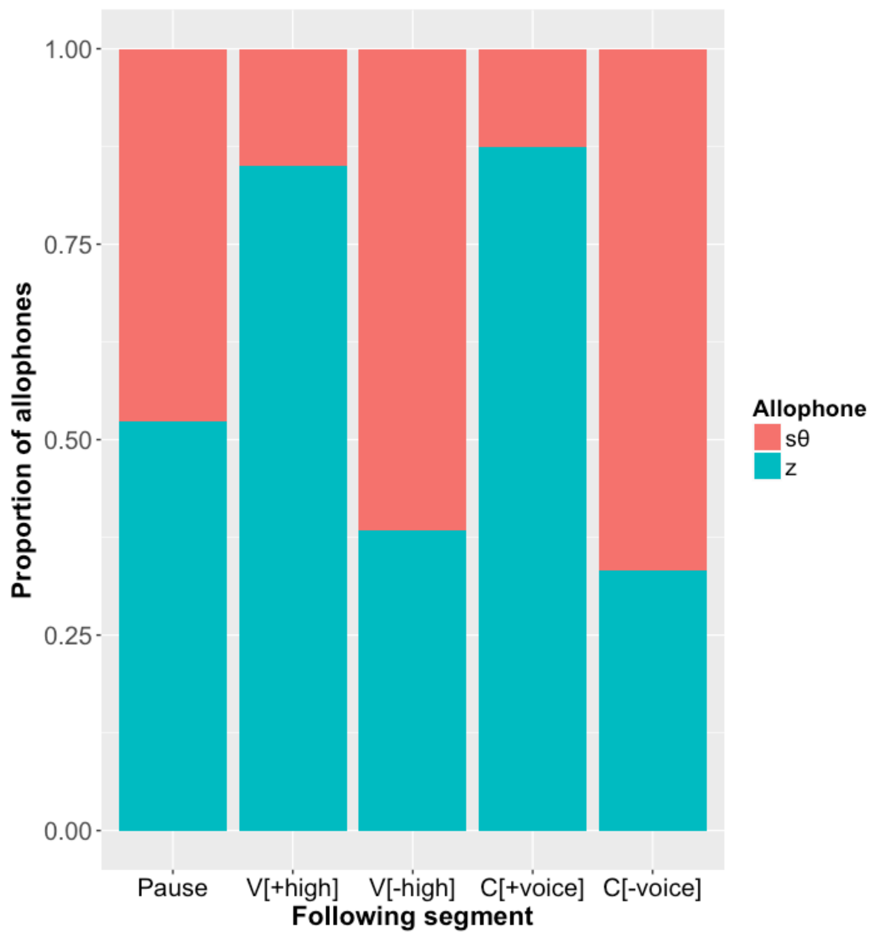

Figure 79. Proportion of [z] and [$s^θ$] produced in coda positions according to following segment.

However, even when the following segment is taken into account, it appears that [z] is still more preferable in coda positions overall: it is strongly preferred before high vowels and voiced consonants (85% and 88%, respectively) while [$s^θ$] is only slightly preferred before non-high vowels and voiceless consonants (62% and 67%, respectively). It is possible, therefore, that this additional effect is due to gender preferences, as [z] is strongly associated with men (see Section 6.5.4) and approximately 60% of the tokens visualized in Figure 79 were produced by male speakers.

With respect to PRESERVE(voice), we saw that additional penalties were greatest in phrase-initial and syllable-final positions, suggesting that the preservation of the input specification [-voice] is sensitive to the salience of the prosodic position in which /s/ occurs. Indeed, Figure 80 shows that the likelihood of preserving voicing specifications decreases dramatically as the prosodic position gets weaker.

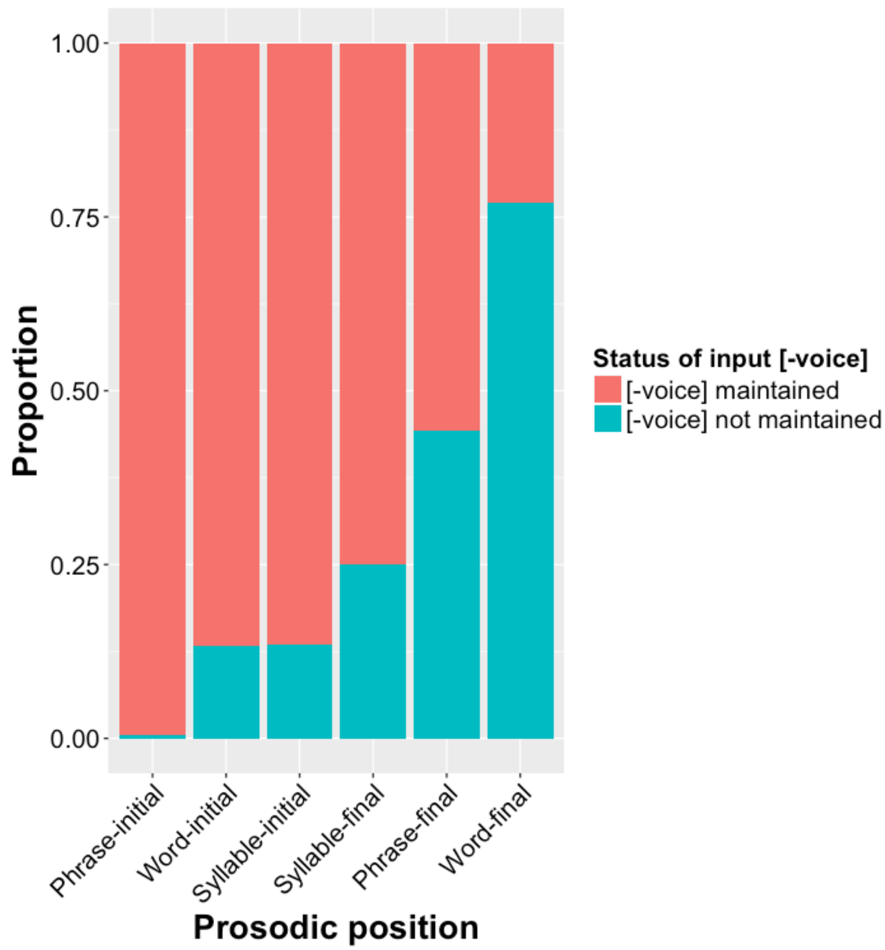

Figure 80. Status of input [-voice] specification by prosodic position.

In phrase-initial position, [-voice] is almost categorically preserved, with [s] and [s$^{\theta}$] representing the most frequent outputs. In syllable-final position, [-voice] is preserved much more often than in word-final position, as [h] is highly preferred to [ɦ] in the former. For example, for input *escuela*

/eskwela/ (‘school’), the grammar assigns predicted probabilities of 35.8% and 3.6% to outputs [ehˈkwe.la] and [eɦˈkwe.la], respectively. In word-final position in the same phonological environment, as in *es que* /eske/ (‘it’s that’), this difference shrinks to 25.9% versus 17.1%.

With respect to PRESERVE(coronal), the largest additional penalties are incurred in word- and syllable-initial positions, as the majority of variation in these strong positions is among [s], [z], and [s^{θ}]. This divide is a particularly important one, as the complete loss of the oral gesture in these positions likely makes prosodic boundaries between words and syllables less perceptible. That is, while the impetus to ease articulatory effort cost is always present, this tendency “is resisted in strong prosodic positions; and as a result, segmental cues may be weaker in weak positions” (Keating, 2006, p. 180). What we see with respect to PRESERVE(coronal) is this resistance to weakening an important segmental cue, [+coronal], in strong prosodic positions. Figure 81 compares allophonic variation for three intervocalic contexts to demonstrate the importance of preserving the [+coronal] cue in more salient prosodic positions as compared to the less salient, word-final position.

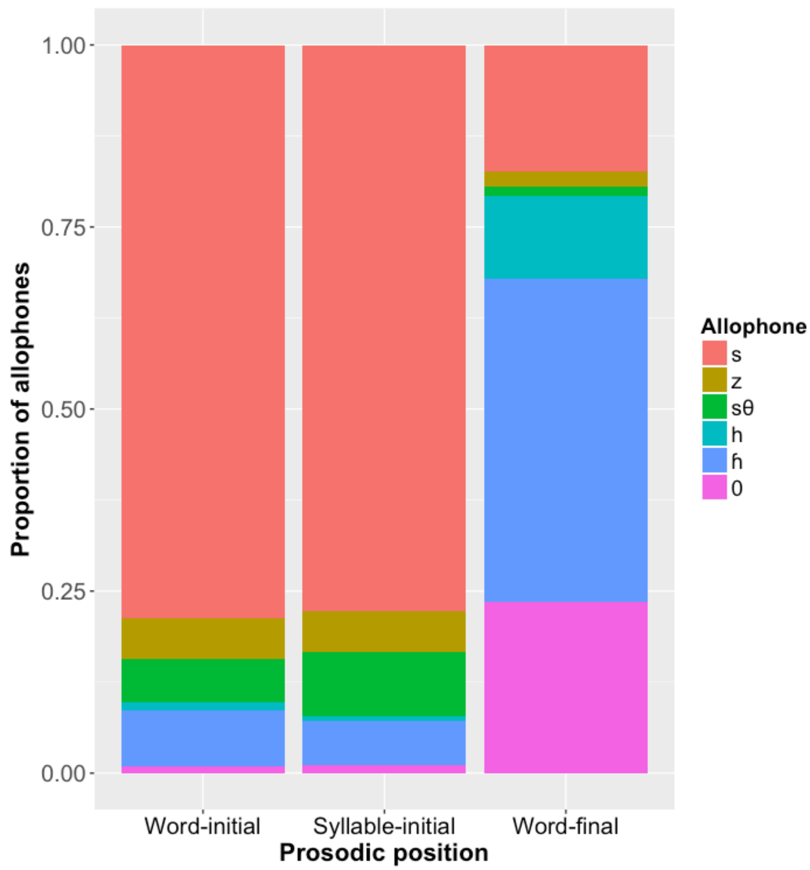

Figure 81. Allophonic variation for intervocalic tokens of /s/ in three prosodic positions.

In Figure 81, the red, yellow, and green segments of each bar represent realizations for which the [+coronal] specification has been preserved. When all else is created equal, word-initial and syllable-initial tokens of intervocalic /s/ preserve coronality at much higher rates, as is reflected by the additional penalties incurred for PRES(cor) in these positions. Because the only difference here is prosodic position, the argument that the observed patterns can be attributed to the weakening of segmental cues in weak prosodic positions is a compelling one.

Finally, PRESERVE(segmental) behaves similarly to PRESERVE(coronal) in that we observe a dramatic drop in additional penalties between syllable-initial and syllable-final positions. Furthermore, while the effects of tonicity were unclear for PRESERVE(strident), PRESERVE(voice), and PRESERVE(coronal), PRESERVE(segmental) behaves similarly to PRESERVE(strident, voice) in that tonicity plays a clear role in determining weights. These results suggest that syllable stress may only be important for determining patterns of /s/ retention vs. lenition and segmental absence vs. presence, but not for intermediate realizations. Because the majority of accounts of Spanish /s/ weakening work within a tripartite system of allophones, it's difficult to situate these findings within the existing literature. However, a number of studies on Caribbean varieties of Spanish including Terrell (1979), Alba (1982), Lafford (1989), Alfarez (2000), and Lynch (2009) find that both word-medial and word-final coda deletion are more likely in atonic syllables, while the sibilant is favored in tonic syllables. It is possible that only [s] and [∅], the most extreme but frequent allophones in the Salvadoran data, are sensitive to syllable stress.

A crucial difference between weights for PRESERVE(segmental) and the other four Faithfulness constraints is that violations in word-final position incur greater penalties than in syllable- or phrase-final positions. While word-final position is comparatively weaker than the other two coda positions, this finding is likely an indicator of hiatus avoidance: when /s/ is deleted in an intervocalic context, as in the noun phrase *los hombres* [loˈom.bres] 'the men,' it creates hiatus: heterosyllabic vowel-vowel sequences, which are dispreferred in Spanish (Hualde & Prieto, 2002; Quilis, 1993). In an effort to mitigate the perceptual cost of potential hiatus, it is likely that higher weights for PRESERVE(segmental) in word-final position serve to deter complete elision of /s/ in certain contexts. Figure 82, below, confirms this hypothesis: while lenition is rampant in word-final position regardless of following context, deletion is much more prevalent preconsonantally.

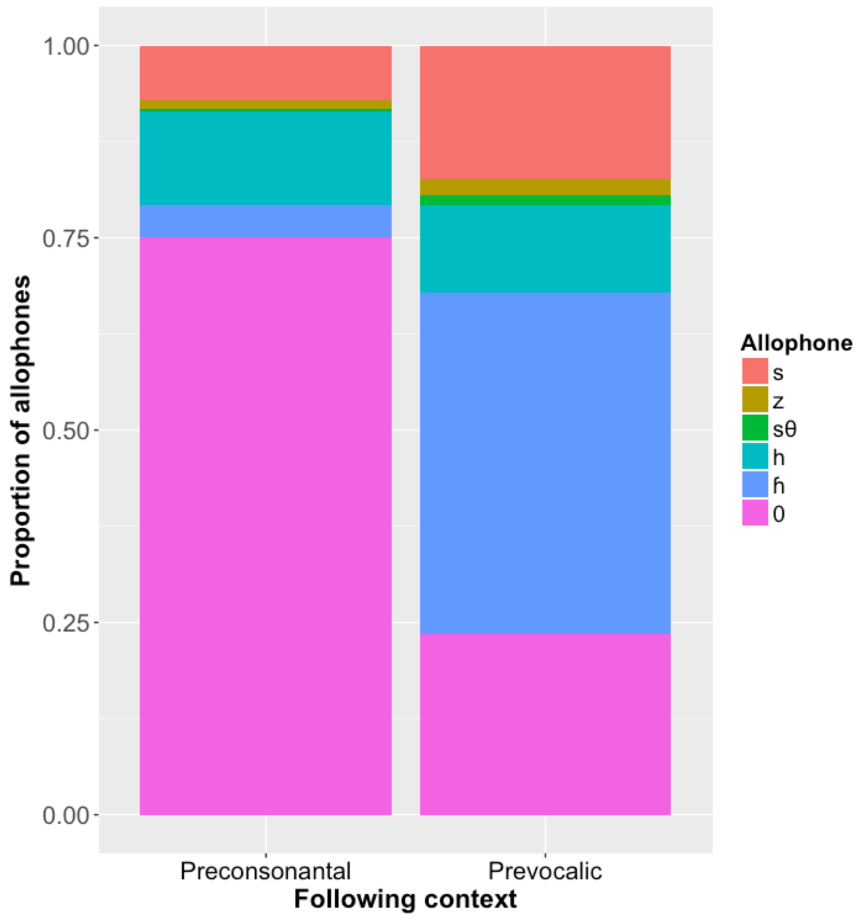

Figure 82. Allophonic variation in word-final position by following context.

In prevocalic position, [ɦ] is the most frequently-occurring allophone and as it is able to serve as a low-effort substitution without compromising important perceptual contrasts.

Also central to this issue of hiatus avoidance is the relationship between /s/ weakening and syllabification. That is, because consonants serve as onsets in prevocalic contexts (including those that occur word-finally in connected speech), a word-final /s/ that occurs in coda position at the word level resyllabifies as an onset when it occurs prevocalically in connected speech (Harris, 1993; Hualde, 1992). For example, compare the word-final /s/ in *unas* [u.nas] 'some' to that in *unas alas* [u.naˈsa.las] 'some wings': in the former, [s] serves as a syllable coda, while in the latter it resyllabifies as an onset in the syllable [sa], with lenition and syllabification interacting across word boundaries (Face, 2002). Therefore if /s/ is expected to resyllabify as an onset in word-final, prevocalic contexts, it is possible that the stress of the following syllable helps determine the perceptual cost of deleting said /s/. Indeed, accounts of /s/ weakening in dialects such as those of Cuban and the Dominican Republic (e.g., Alba, 1982; Dohotaru, 1998; Poplack, 1979; Terrell, 1979) find that rates of word-final deletion drop precipitously before stressed vowels. Figure 83 confirms the same pattern in the Salvadoran data.

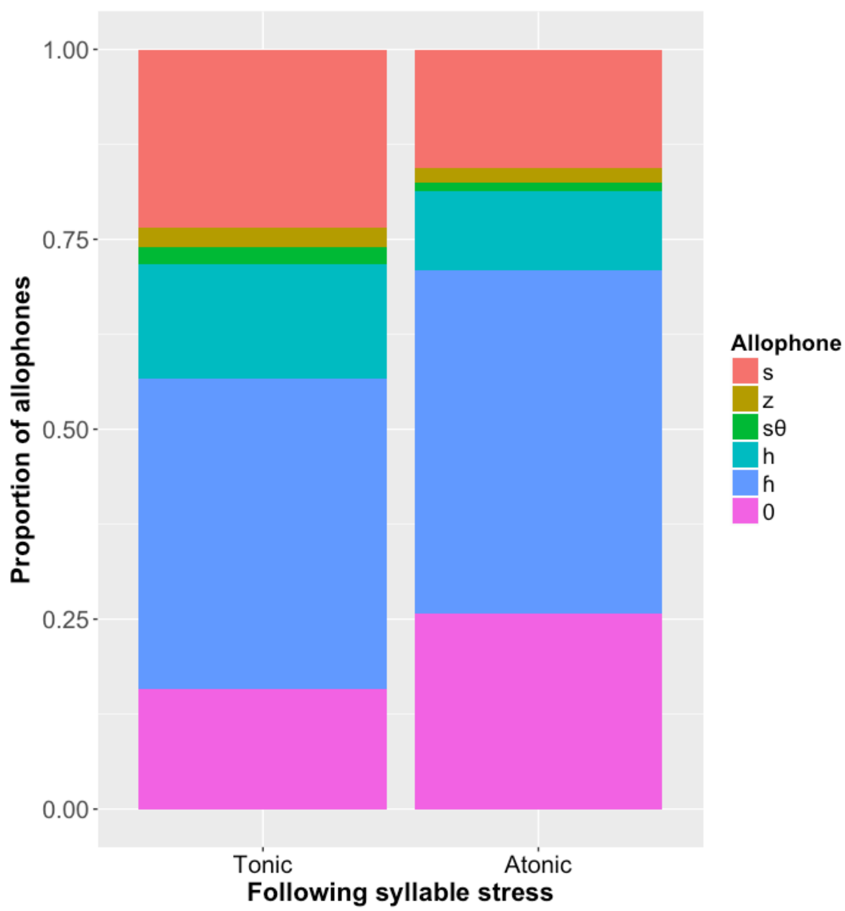

Figure 83. Allophonic variation in word-final position in prevocalic contexts according to the tonicity of the following syllable.

Consistent with previous findings, we see that deletion rates drop significantly in word-final position in prevocalic contexts when the following vowel is tonic. It is likely that these asymmetries are due to a cross-linguistic peculiarity: stressed, word-initial vowels are uniquely resistant to elision as a strategy for hiatus resolution (Casali, 1998), thereby increasing the need for an intervening segment. However, while other authors report high rates of retention of [s] in these contexts, Salvadoran speakers are most likely to utilize intermediate weakened allophones in this weak prosodic position.

7.3 Socially-conditioned variation

While the base grammar reflects fundamental patterns in the /s/ weakening data according to language-internal factors, Section 6.5 presented changes in the relative importance of Faithfulness when constraints are modulated according to a speaker's region of origin, urban vs. rural place of origin, age, gender, and education. Table 60 summarizes these findings.

Table 60

Relative importance of Faithfulness for social groups compared to the base grammar

Group	Average Markedness constraint weight	Average Faithfulness constraint weight	Relative importance of Faithfulness	Scaling factor
El Salvador (all)	1.29	2.16	0.87	N/A
Region				
Santa Ana	1.05	2.00	0.95	+0.08
San Miguel	1.41	2.13	0.72	-0.15
San Salvador	1.21	2.00	0.79	-0.08
Origin				
Rural	1.37	2.13	0.76	-0.11
Urban	1.15	2.04	0.89	+0.02
Age				
Older	1.32	2.04	0.72	-0.15
Younger	1.21	2.24	1.03	+0.16
Gender				
Women	1.50	2.21	0.71	-0.16
Men	1.03	2.06	1.03	+0.16
Level of education				
Lower	1.17	2.04	0.86	-0.01
Higher	1.07	2.08	1.01	+0.14

Evident here is that San Migueleños, San Salvadoreños, rural speakers, older speakers, women, and those with lower levels of education assign more importance to reducing effort cost and less to preserving perceptual cues for /s/, particularly in more salient prosodic positions. On the other hand, speakers from Santa Ana, from urban communities, those who pertain to the younger age group, men, and those with higher levels of education prioritize the preservation of perceptual cues more than the average Salvadoran. Differences are greatest with respect to age and gender—two factors that are known to condition /s/ lenition in other dialects of Spanish—but not in the directions we might expect based on the literature.

The following sections provide brief discussions of findings related to the above social groups. In particular, these sections strive to highlight findings that mirror those of other /s/ weakening dialects and propose possible explanations for those that deviate.

7.3.1 Region

Findings for variation by region confirm an important distinction posited by Azcúnaga Lopez's (2010) linguistic atlas of El Salvador, which characterizes the San Miguel region as the most phonologically nonstandard. My results show that San Migueleños prioritize Faithfulness much less than their counterparts in Santa Ana and San Salvador, consequently leniting /s/ at higher rates, utilizing more allophonic variation outside the traditional tripartite system, and doing so in more salient prosodic positions. While Azcúnaga Lopez divides the country into only two linguistic regions (the occident—encompassing both Santa Ana and San Salvador—and the orient), I hypothesized that speakers from Santa Ana would show evidence of standardization due to their close proximity to Guatemala, whose phonology "resembles that of Costa Rica more than that of neighboring El Salvador" in part because "Guatemalans pronounce /s/ with clarity as a tense, grooved sibilant" (Canfield, 1981, p. 55). While it is difficult to confirm that dialect contact is the source of standardization, speakers from Santa Ana do retain /s/ at the highest rates despite reporting the lowest levels of education among the three groups (see

Figure 72). Furthermore, many participants from Santa Ana spoke of frequent travel to Guatemala, and many had friends or family members who had relocated there for economic reasons or to escape the pervasive gang violence of El Salvador. Anecdotal accounts of Salvadoran migration to Guatemala, which might result in the formation of transnational families and possible dialect leveling by osmosis, are confirmed by Pizarro (2002), who writes that this movement increased during the Salvadoran civil war of the 1980s and again in the early 2000s. However, it's difficult to identify the regions of origin of these immigrants in the census data.

7.3.2 Origin

Findings with respect to urban versus rural origin confirm findings in both Spain and Latin America that urban speakers, who are exposed to more standardized forms of language, lenite /s/ less than rural speakers (see Chapter 2 for a review of these studies). Furthermore, scholars such as Canfield (1953), Lipski (1994), and Hualde (2005) have posited that the use of the allophone I refer to as Salvadoran [s^{θ}] is distinctly rural. This hypothesis was empirically confirmed, although it should be noted that use of this allophone is not limited to rural speech but rather is a characteristic of the speech of women, the older generation, and San Migueleños, all of whom were balanced for origin.

While Cedergren (1973) finds that rural speakers in Panama tend to weaken /s/ to [h] while urban speakers are more likely to delete /s/ entirely, these differences are not obtained in the Salvadoran data. Rural speakers are, in fact, more likely than their urban counterparts to produce all non-[s] allophones, including but not limited to [h], [ɦ], and [∅]. Additionally, while the urban/rural divide is thought to be extremely salient in Spanish dialectology, the distinction between urban and rural speakers was the smallest of any social group comparison presented in Chapter 6. This blurring of urban and rural speech communities in El Salvador could be due, in part, to a mass migration to urban areas that began in the wake of the civil war and persists as intense gang violence continues to plague rural areas today. Indeed, The World Bank's data catalogue ("Urban population (% of total) | Data," 2018) chronicles changes in the percentage of Salvadorans who live in urban communities, showing a sharp increase from less than 40% in 1971 to approximately 67% in 2016. While I was strict with my demographic selection criteria,[37] it is

[37] I only considered city-dwelling participants 'urban' if they were born in the city in which they currently lived and had not resided elsewhere for more than six months. Participants were only considered 'rural' if they were born in the rural community in which they currently lived and had not resided elsewhere—including to attend university—for more than six months. Prospective participants who did not meet these criteria were not recruited.

possible that increased movement between rural and urban areas over the past 50 years has resulted in sociolect leveling.

7.3.3 Age

While younger speakers tend to be leaders in linguistic change, this can mean one of two things for /s/ lenition: either they innovate away from the linguistic standard and weaken /s/ at higher rates than their older counterparts (e.g., Alba, 1990; Cedergren, 1973), or they prefer the most prestigious variant (usually, but not always, [s]) to avoid association with dispreferred demographics (e.g., Carvalho, 2006; Lafford, 1986; Lynch, 2009). The situation in El Salvador appears to be a case of 'change from above' like that observed by Lafford in Colombia: younger Salvadorans are, on average, better educated and more literate than older generations ("El Salvador National Education Profile," 2014) and consequently speak more prescriptively.

If the move toward more standard phonology is, indeed, a change from above led by younger Salvadorans, the implications for the future of Salvadoran Spanish are paramount. As I discuss in Chapter 1, El Salvador's historical isolation has played an important role in shaping its dialect (Lipski, 1994). That is, by providing a barrier to linguistic standardization, an isolated El Salvador has been able to both preserve archaic forms of language from the colonial period and innovate away from an international prestige. In an increasingly globalized and technologized world, however, Salvadorans are suddenly more connected than ever,[38] and an apparent-time interpretation of the /s/ weakening data would propose that differences between older and younger

---

[38] The Central Intelligence Agency ("The World Factbook," 2018) reports that internet usage in El Salvador grew almost 400% between 2007 and 2015, and that there were approximately 148 mobile phone subscriptions per 100 residents in 2016.

speakers reflect this shift. Of course, it is possible that age-based differences obtained in this study do not reflect sound change in progress but rather age-graded variation, which supposes that speakers modify their linguistic behavior over the course of their lifetimes. However, particular differences in the speech of older and younger Salvadorans such as the prevalence of [s^{θ}]—which by all accounts originated in precolonial Andalusia—in the former are more difficult to justify via this theoretical approach. These issues will be revisited in Section 7.4, which addresses the need for longitudinal linguistic research in El Salvador.

7.3.4 Gender

Social factor comparisons according to region, origin, and age have painted a fairly clear picture in which the group that prioritizes Markedness produces less tokens of [s] and more tokens of each of the five lenited allophones in its place, while the group prioritizing Faithfulness does just the opposite. Findings with respect to gender, however, are more nuanced: women produce more of the variants [s^{θ}], [h], and [ɦ], while men produce more [s], [z] and [∅]. While overall rates of lenition are not strikingly different between the two groups, the variants used more frequently by women are those that carry the most social stigma, as attested both in sociolinguistic interviews[39] and in previous research describing speaker evaluations of these sounds as rural and uneducated. Non-standard variants dominated by men, on the other hand, are likely social markers but not necessarily stereotypes. That is, use of these variants depends in part on social group and speech style, but speakers do not comment on them overtly and they show less pronounced social stratification (Labov, 1972). These results not only contradict previous research with respect to

[39] In sociolinguistic interviews, participants frequently commented on the use of "la jota" and "la zeta" (i.e., [h] and [s^{θ}]) when asked about the speech of Salvadorans from other regions, communities, or socioeconomic groups. Language attitudes with respect to these variants revealed pervasive stereotyping, including by those who produced these variants at very high rates themselves.

gender and /s/ weakening in Spanish but are difficult to explain within greater sociolinguistic theory. This section will explore these issues in depth as they may have important implications for the status of /s/ weakening in Salvadoran Spanish.

The accepted creed in sociolinguistics with respect to language and gender stems from Labov's (2001) gender paradox, which asserts that, when a linguistic variable is stable, women utilize more prestigious variants compared to men of the same age and social class. Furthermore, when a linguistic change is in progress and this change is driven by overt prestige, women are the quickest to adopt it; however, women are also more likely than men to use innovative (non-prestigious) linguistic forms when change is driven by covert prestige. This theory has been tried and tested across cultures and languages, and these principles are well-equipped to explain gender-graded patterns of /s/ weakening in other dialects of Spanish (see Chapter 2 for a review of this literature). In the Salvadoran data, however, it appears that none of these principles hold sway. Let's first reexamine the gender-graded patterns presented in the previous chapter, this time by age and education for the most accurate comparison. Figures 84a and b show allophonic variation by gender and age for the lower and higher educational groups, respectively.

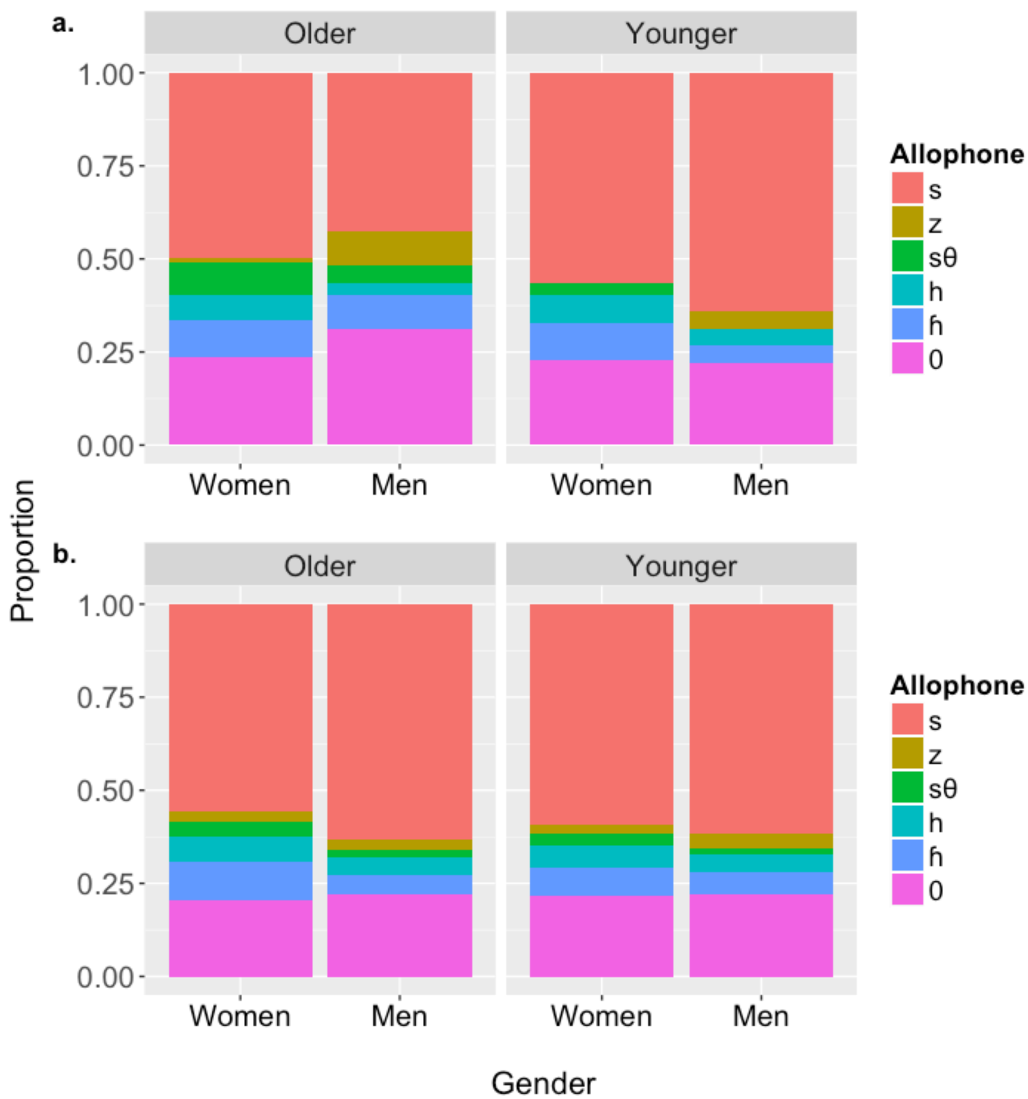

Figure 84. Allophonic variation by gender and age for lower (a.) and higher (b.) educational groups.

Here we see that women produce stigmatized variants at higher rates than men regardless of age or level of education. Even older male speakers with less education, who produce less [s] than their female counterparts, still produce [s^{θ}], [h], and [ɦ] at lower rates. Given the comparisons presented here, Labovian principles indicate that the /s/ weakening variable is not stable. However, if Salvadoran /s/ lenition is undergoing a change toward standardization—as is suggested by age-based differences, literacy rates, and increased access to education and technology—this change from the top should be led by women, which is clearly not the case.

One possible explanation for these unexpected results is that sociocultural, political, and economic forces unique to El Salvador have shaped gender differences in such a way that men are poised to lead the change from above. Slow economic growth in El Salvador has resulted high unemployment that disproportionately affects women, who are less likely to be literate and complete fewer years of schooling on average than their male counterparts. As a result, the ratio of female to male labor force participation is strikingly lower than in other low, middle, and high-income countries ("Gender Data Portal | Country - El Salvador," 2018). As reported in sociolinguistic interviews, this lack of economic opportunity combined with a fear of gang violence means that many Salvadoran women seldom leave the house, effectively isolating them from their communities. The role of women in leading change from above has been attributed to a desire to avoid social judgement (Gordon, 1997; Trudgill, 1972) as well as the use of symbolic capital such as language and clothing to signal social group membership (Eckert, 2000). However, if women in El Salvador aren't participating in society in the way conceived of by these theories, these interpretations may not be valid.

Indeed, sociolinguistic work on Arabic—a language spoken in places in which women tend to play a significantly smaller role in public life than their male counterparts—reveals a relationship between language and gender that is similar to that observed in the Salvadoran data. That is, in diverse communities from Cairo (Haeri, 1987) to Morocco (Sadiqui, 2003) to Saudi Arabia (Al-Essa, 2008, 2009), women use less standard linguistic forms than men, marking a "widespread reversal of the positions of men and women predicted" by the Labovian gender paradox (Labov, 2001, p. 270). This gender dynamic has been a topic of interest for Arabic linguists for some time, beginning with variationist interpretations that Arab women have little opportunity to access standard linguistic forms through venues such as education and literacy (Al-Wer, 2014). Scholars outside the variationist tradition have taken a similar but more nuanced approach to this question, arguing for what Sadiqui (2003) calls "sociolinguistic androcentricity": women are excluded from the "sites of power" around which society revolves, as men dominate the important social institutions of religion, politics, and law. From this perspective, Al Wer (2014) argues that we must examine the factors specific to a culture and individual community that give rise to labels such as "prestige" and "stigma." Within this framework, it is possible that the unique isolation of Salvadoran women—particularly those in the older generation—has created a case of gender-based sociolinguistic segregation similar to that observed in the Arab world.

Increased movement of men as compared to women between El Salvador and the United States, reported by many of my participants,[40] could also have implications for linguistic change. The United States is a site of rich dialect contact, and many Salvadorans in the U.S. interact intimately with speakers of Mexican Spanish and other more prestigious /s/-retaining dialects at work and in the community, resulting in dialectal accommodation over time (Parodi, 2005, 2011). If these speakers return to El Salvador[41] or maintain transnational relationships with friends and family, these dialectal changes could permeate particular speech communities in the home country as men lead change from above.

[40] Given the challenge of accounting for undocumented individuals in official documentation, these observations are difficult to corroborate, but in 2008, of all the Salvadoran foreign-born legal permanent residents in the United States, over 60% were men (Terrazas, 2009). Additionally, while the gender ratio for 15-24-year old Salvadorans is 1.01 men to every woman, this number drops to 0.86 for the ages of 25-54 despite an average life expectancy of 74.9 years, suggesting that men of prime working age are leaving the country at high rates ("The World Factbook," 2018).

[41] A 2014 report by U.S. Immigration and Customs Enforcement (ICE) obtained by the Transactional Records Access Clearinghouse (TRAC) at Syracuse University under the Freedom of Information Act (FOIA) shows that more than 21,000 Salvadorans were deported from the U.S. during fiscal years 2012 and 2013; more than 90% of them were male.

In addition to raising questions about the status of the /s/ weakening variable in Salvadoran Spanish, the data highlight a clear gender divide with respect to [z] and [s$^{\theta}$]. It appears that these allophones, which are phonologically similar in that they prefer the same contexts (intervocalic, C__V) and prosodic positions (word-initial, syllable-initial) and only differ from [s] by a single feature, are different means to the same end: both serve to reduce effort cost while preserving important perceptual distinctions in strong prosodic positions, but [z] is strongly preferred by men and [s$^{\theta}$] by women. Variable voicing of /s/ has, in fact, been associated with men in other dialects of Spanish including those of Ecuador (García, 2015), Spain (McKinnon, 2012), and Costa Rica (Chappell, 2016). In fact, in a study that investigates social perceptions of intervocalic voicing of /s/ in Costa Rica, Chappell (2016) shows that men who produce [z] intervocalically are perceived to be significantly nicer, more confident, more Costa Rican, and more masculine than those who only produce [s], while there is no effect for female speakers. Both genders, however, are perceived as pertaining to a lower socioeconomic group when using [z]. Faced with these differences, Chappell argues that men are able to access the covert prestige of intervocalic voicing and therefore benefit from using it despite negative socioeconomic associations; women, contrastingly, don't stand to gain anything from use of the nonstandard variant and are therefore less likely to produce it.

Chappell's interpretation, situated within third-wave sociolinguistic theory (Eckert, 2012), might also be applied to the gendered use of [s$^{\theta}$] in El Salvador. While this sound isn't articulatorily interdental, the fronting of the tongue body gives it an acoustic quality similar to that of an interdental sound and it has been described as approaching a 'lisp' (see Chapter 4 for a thorough account of this sound and its acoustic/articulatory properties). While there is no scientific evidence to support the stereotype that lisping is effeminate or a feature of gay speech, it is nevertheless a pervasive stereotype. In a matched-guise experiment on American English, Munson and Zimmerman (2006) found that listeners were significantly more likely to rate talkers as "gay-sounding" when /s/ was "frontly misarticulated," which they describe as a quasi-interdental sound with a COG of 3068 Hz, a measure very similar to the one found for [$s^θ$]. This was compared to a "neutral /s/" with a COG of 6677 Hz, a measure very similar to the one found for Salvadoran [s]. The authors conclude that the 'gay lisp' stereotype is incredibly powerful. While a study of this nature is yet to be done in Spanish, Mack's (2010, p. 140) study of gay speech stereotypes in Puerto Rico reveals that the most salient phonetic cue implicated in the speech stereotype of gay men was "distinct pronunciation of /s," while the most salient cue overall was "sounds feminine." It is possible, then, that the acoustic qualities of [$s^θ$] in El Salvador convey characteristics of femininity, sensitivity, etc. which, despite also indexing ruralness or low socioeconomic status, have a social benefit for women yet pose an additional social cost for men in this hyper-*machista* society. While this interpretation is speculative at best and would require corroboration via targeted matched-guise experiments in Spanish, the high degree of contact between English and Spanish makes this line of inquiry worth pursuing further.

7.3.5 Social group interactions

Examining social group interactions revealed patterns that were obscured when looking at each social factor independently. First, the interaction between region and age revealed that significant age-based differences in variant production affect different sounds in different regions. In San Miguel, the largest difference is with respect to [$s^θ$]; in San Salvador, [z]; and in Santa Ana, [∅]. That is, younger speakers in each of these regions are much less likely to produce the given variant than the older generation. With respect to [$s^θ$] in San Miguel this effect is likely an artifact of a general shift toward standardization, as this variant is most prevalent in this region overall and therefore most susceptible to change from above. With respect to [z] in San Salvador, this effect may reflect the general tendency in the capital to minimize use of variants outside the traditional tripartite system (while San Salvadorans lenite /s/ at higher rates than speakers from Santa Ana, they produce less [s^θ] and [z] overall). Finally, in Santa Ana, the dramatic drop in /s/ deletion by younger speakers could be indicative of a more dramatic standardization effect in which even the most ubiquitous patterns of /s/ weakening are changing.

An examination of gender differences by region further highlights Santa Ana as a unique social environment. While the gender differences discussed in the previous section are extremely prevalent in San Salvador and San Miguel, women in Santa Ana are in fact significantly less likely to produce all nonstandard variants except [h] and [ɦ], for which men and women are not significantly different from one another. Given these gender discrepancies and the more dramatic trend toward standardization in this region, the Sana Ana data aligns well with Labovian principles: it appears that women are leading a change from above. Crucially, the interaction of these two factors reveals that the perplexing gender patterns described in the previous section are constrained by region. It's difficult to posit why women in Santa Ana pattern so differently from their counterparts in San Salvador and San Miguel without further research into this line of inquiry, particularly because the most recent census data available indicates that unemployment rates are higher in Santa Ana than in the other two regions and that speakers in San Salvador have the highest literacy rates of the three groups (Gobierno de la República de El Salvador Ministerio de Economía Dirección General de Estadística y Censos, 2007).

Finally, interaction effects for gender by origin revealed that the urban/rural distinction is much more important for women than men. That is, with the exception of [z], differences in variant production for urban and rural women were always larger than for urban and rural men. This is likely due to the exacerbation of gender inequality with respect to (socio)economic opportunities in less-developed, rural areas. Interestingly, while data examining allophonic variation by just origin shows that rural speakers are more likely to produce all nonstandard variants of /s/ than urban speakers, this does not hold true when broken down by gender. For example, we see that while rural men are more likely than urban men to produce [z], urban women are more likely than rural women to use this variant. It is possible that, because [z] indexes masculinity, there is a potential benefit for its use by a population of women more immersed in the workforce. Conversely, we see that while urban men are more likely to produce [ɦ] than rural men, rural women prefer this variant as compared to urban women.

This section and the previous have summarized important features of the maxent grammar as well as patterns of socially-conditioned variation. Even more importantly, these sections have explored the implications of findings by drawing connections to previous work on /s/ lenition, phonetic and phonological theory, and social/cultural/economic conditions specific to El Salvador. The next section compares these findings with John Lipski's work on Salvadoran /s/ weakening from the 1980s in an effort to create a more longitudinal picture of this phenomenon.

7.4 Salvadoran /s/ weakening: Then and now

As I discuss in Chapter 1, empirical research on the Spanish of El Salvador (and on Central American dialects in general) has been neglected within the field of Hispanic linguistics. Despite this gap in the literature, this section attempts to compare findings on Salvadoran /s/ weakening from the 1980s with those presented here in order to address the following questions: do patterns observed in this dissertation align with those documented by Lipski? If not, is it reasonable to make inferences about longitudinal change, or is the available data not sufficient? In order to address the first question, Table 61 compares production rates for allophones of /s/ within the traditional tripartite system with those presented by Lipski (1984). Because Lipski does not differentiate between [h] and [ɦ], I have collapsed data for these two allophones for the purpose of this comparison. The reader should also note that percentages for Brogan (2018) do not add up to 100% as tokens of [s^{θ}] and [z] have been omitted.

Table 61

Comparison of allophonic variation from Lipski (1984) and the present study

Context	Lipski (1984)			Brogan (2018)		
	[s]	[h]	[∅]	[s]	[h]	[∅]
sC	55%	44%	1%	62%	14%	23%
s#C	10%	71%	19%	7%	16%	73%
s##	86%	12%	2%	51%	5%	42%
s#$V_{[+stress]}$	44%	47%	9%	24%	56%	16%
s#$V_{[-stress]}$	28%	69%	3%	16%	56%	26%
V#s$V_{[+stress]}$	99%	1%	0%	85%	5%	0%
V#s$V_{[-stress]}$	91%	9%	0%	75%	11%	0%
Vs$V_{[+stress]}$	95%	5%	0%	85%	1%	1%
Vs$V_{[-stress]}$	89%	11%	0%	75%	9%	1%

Before addressing the issue of whether this comparison is even valuable, let's first assume that it is for the sake of highlighting similarities and differences in the 1984 and 2018 data. Overall, the comparisons presented in Table 61 chronicle a striking move toward lenition. Preconsonantally, where lenition rates were already high in 1984, the 2018 Salvadoran data show a clear shift from [h] to [∅] regardless of prosodic position. Similarly, we see shift to deletion in phrase-final position; in fact, my data show higher rates of lenition phrase-finally than in syllable-final, word-medial position while Lipski's show the opposite. With respect to prevocalic contexts, the effect of stress in the 1984 and 2018 data is extremely consistent: the tonicity of the vowel following /s/ in all prosodic positions. Otherwise, however, we continue to see a dramatic shift toward [h] and [∅], even in word-initial and word-medial onset positions. The comparison in these last two positions is, of course, flawed because 10-15% of realizations in the 2018 data are comprised of [s^{θ}] and [z].

While we would like these comparisons to be valuable, it's important to recognize significant differences in methodologies and analyses between the two studies. The first serious issue is that the data source is unknown for Lipski (1984).[42] However, a study published the following year presents Salvadoran data that is almost identical; it is likely that the data from Lipski (1985), based on recordings of 10 speakers from San Salvador, is also the source of the 1984 data. If this is the case, the above comparison is difficult to justify, as we have seen in the present study that region—in addition to other non-linguistic factors—is incredibly important in understanding the observed variation. Furthermore, the 1984 data were analyzed impressionistically while the present study relied on the waveform and spectrogram in *Praat.* Impressionistic segmental transcription can be extremely difficult, particularly when distinguishing between allophones such as [h] and [∅]. For this reason, it's possible that some of the apparent longitudinal changes simply reflect access to more precise analytical tools and technologies, not dramatic shifts toward lenition over the past 30 years.

In the interest of exploring longitudinal trends in Salvadoran /s/ weakening, this section has presented a comparison with data from Lipski (1984). However, longitudinal trend studies are only valid if the data come from different but *comparable* individuals (Wagner, 2012), so I

[42] In this paper, a footnote indicates that the data is from a forthcoming paper titled "Central American Spanish in the U.S.: the case of El Salvador". However, it appears this paper was never published and I am therefore unable to report the number of speakers or any other demographic information with certainty.

encourage the reader to interpret Table 61 with caution. I will return to the important issue of longitudinal linguistic research in Section 7.8.

7.5 Evidence against pathways of diffusion

In this section, I revisit two interrelated questions of immense interest to scholars in the field, the perspectives for which are presented in Section 2.5:

1. What are the pathways of diffusion for /s/ weakening? (i.e., where does /s/ lenition start, and how does it progress over time?)
2. Are onset and coda /s/ weakening one phonological variable or two?

In Chapter 2, I argue that, within an OT framework, it is not necessary to consider pathways of diffusion because patterns are determined by competing constraints in the grammar. However, now that I have presented my full analysis of the Salvadoran data, it is important to briefly address how previously-proposed theories are not able to account for observed patterns. In order to do so, it's helpful to first examine rates of lenition in the prosodic positions and phonological contexts that previous proposals have relied on in their arguments. Figure 85 shows rates of lenition in the following environments: word-final preconsonantal, word-final prevocalic, word-medial preconsonantal, word-medial intervocalic, and word-initial intervocalic.

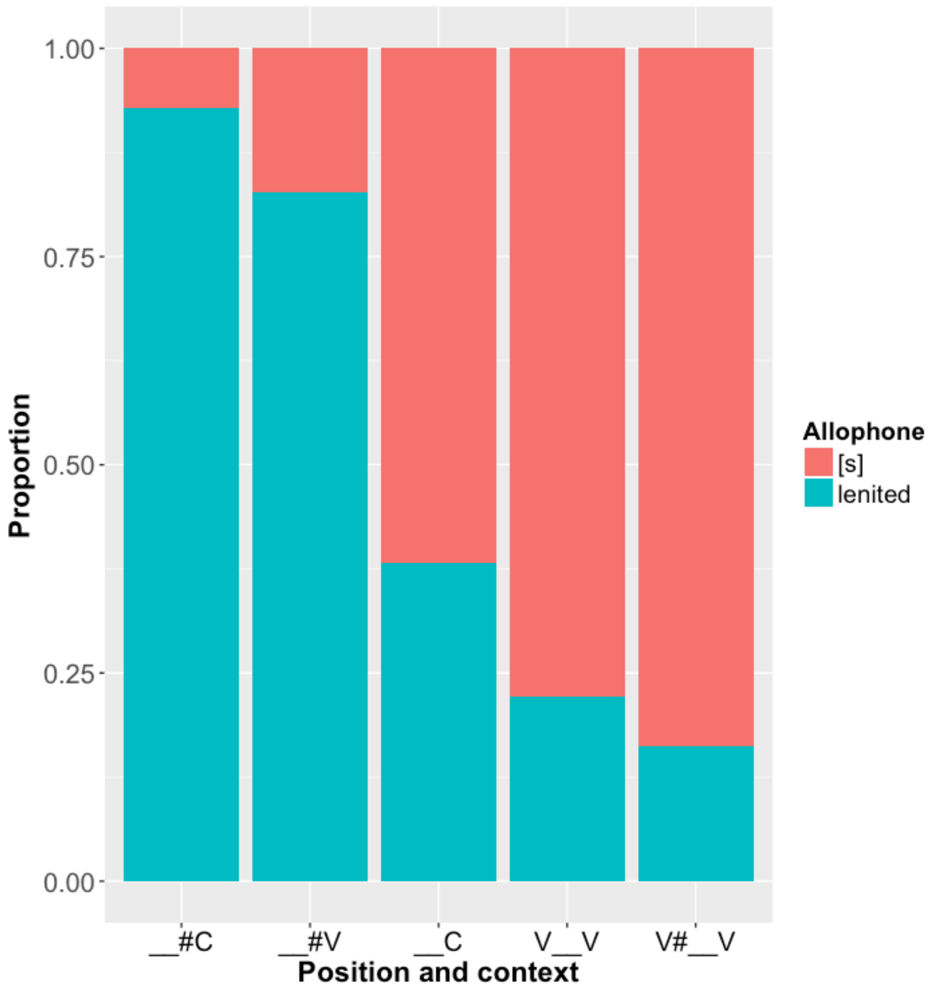

Figure 85. Rates of lenition in the five positional + contextual combinations of interest.

The first proposal to address posits that /s/ weakening follows two distinct pathways: one that begins in syllable-final, preconsonantal position and another that begins in word-medial, intervocalic contexts and then extends to word-initial position. This proposal cannot account for observed patterns in Figure 85, as lenition is much more prevalent in word-final, prevocalic contexts than in word-medial coda position. The second proposal posits that onset and coda /s/ weakening are one process that begins preconsonantally and advances as far as word-final, prevocalic contexts in some dialects; it is only in these dialects that prevocalic /s/ weakening may then be generalized first to word-initial and then to word-medial onset positions. While it is true that the Salvadoran data show evidence of pervasive weakening in word-final, prevocalic position as well as in word-medial and word-initial onset positions, this proposal runs into the same issue as the former: word-final, prevocalic weakening far exceeds word-medial, preconsonantal weakening. Furthermore, Lipski (1999) proposes that word-initial weakening precedes word-medial onset weakening, but this is not the case in El Salvador.

While both of these proposals identify following context as the most important factor for /s/ lenition, it is in fact the prosodic position. In Figure 85, we see that, when we reduce Salvadoran /s/ lenition patterns to these five environments, prosodic position emerges as the most important factor and is then modulated by following context. This is consistent with the premise of this dissertation, which argues that the strength of the prosodic position determines how acceptable lenition is given the need to preserve perceptual cues, while the effort cost imposed by a given context determines the need to lenite; it is these competing factors that define patterns of Salvadoran /s/ lenition. Figure 86 visualizes these factors as tiers.

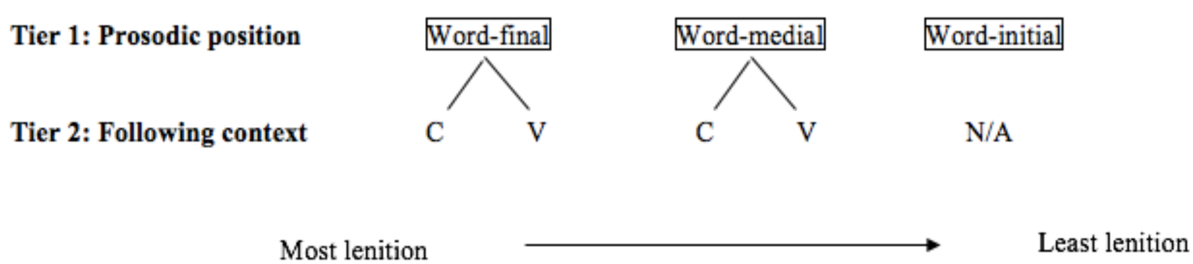

Figure 86. Tiers for Salvadoran /s/ lenition by position and following context.

With respect to the relationship between onset and coda /s/ weakening, Esther Brown and Torres Cacoullos (2003, p. 35) propose that “syllable initial /s/ reduction will occur in Spanish varieties where: (1) word-final /s/ reduction is greater before a pause and a vowel than before a consonant, and (2) overall /s/ reduction rates are relatively low." As we have seen throughout this dissertation, neither of these statements applies to the Salvadoran data. In response to this proposal, Earl Brown and Esther Brown (2012) present a cross-dialectal comparison of /s/ weakening in seven varieties of Spanish, showing that the three dialects with both onset and coda weakening show no significant differences between word-final weakening in preconsonantal and prevocalic contexts, while dialects with only coda weakening do. The authors therefore conclude that a better diagnostic for determining whether a dialect will weaken /s/ in onset position is to compare rates of lenition in word-final contexts. However, this proposal also fails to account for the Salvadoran data, as /s/ lenition in word-final position is, in fact, significantly different preconsonantally and prevocalically ($p<0.001$). The logistic model is presented in Table 62.

Table 62

Logistic model with following context predicting [s] retention in word-final position

Context	Estimate	Standard error	*p* value
Intercept (Preconsonantal)	-2.56	0.08	<0.001***
Prevocalic	1.00	0.11	<0.001***

As we see here, /s/ is significantly more likely to weaken in word-final, preconsonantal contexts than in prevocalic contexts.

While I agree with Earl Brown and Esther Brown's (2012, p. 104) conclusion that "[syllable-]initial /s/ reduction (word-initial and word-medial) … contrasts significantly with intervocalic /s/ reduction in word-final position", I don't believe this is because onset and coda weakening are distinct phonological variables. Instead, this dissertation has argued that differences can be accounted for according to constraints requiring that certain phonetic cues be preserved in salient prosodic positions. While the importance of these constraints relative to those necessitating articulatory effort reduction undoubtedly varies across dialects, sociolects, registers, speech rates, and countless other factors, I believe this framework is better equipped than previous proposals to explain asymmetries between onset and coda /s/ weakening in dialects that show evidence of both.

7.6 Lexical frequency effects

A core hypothesis of this dissertation is that speakers undershoot articulatory targets of /s/ to reduce effort but do so only to the extent that communication is not compromised and target words are still recoverable. Lexical frequency seems indispensable to this theory, as higher frequency words should be more familiar, predictable and have more diverse exemplar representations. As Keating (2006, p. 181) explains, “we would expect to find a relation between predictability (on some independent measure) and degree of strengthening, with greater predictability associated with less strengthening.” In this vein, I hypothesized that a lexical frequency effect would be particularly evident in initial prosodic positions where segments are less predictable because they are less determined by prior context. However, data exploration and statistical testing in Section 6.2 revealed that this factor does not significantly condition patterns of Salvadoran /s/ lenition, even when prosodic position is accounted for. While words containing [s] *did* have lexical frequencies that were significantly lower than words containing nonstandard variants, frequencies for words containing [∅] were the most similar (i.e., had the second-lowest lexical frequencies on average), suggesting that frequency did not have a legitimate effect on allophonic variation. For this reason, lexical frequency was not incorporated into the maxent grammar.

One possible explanation for the lack of a frequency effect is that some other measure(s) of word predictability is at work. It may be, for example, that Salvadoran /s/ lenition is conditioned by lemma frequency instead of (or in tandem with) word frequency such that the predictability for *tenía* ('I had') depends on the frequency of the infinitive *tener* ('to have'), or perhaps on both entries, as they're well-correlated. It is also possible that frequency measures are less important than the semantic predictability or phonological neighborhood density (i.e., the number of other words (or near-words) that could be generated by replacing /s/ with another phoneme) of the target word.

With respect to word-initial weakening, which should be particularly sensitive to frequency as described above, another possible explanation is that word-initial /s/ reduction in a few high-frequency words has been lexicalized, or near-lexicalized, for some speakers. If these words are disproportionately representative of word-initial lenition in the data, wide-ranging frequency effects might not be obtained. Lipski (1986) finds an effect of this nature in neighboring Honduras, where ten high-frequency words beginning with /s/ are disproportionately likely to be produced with a weakened [h]. In order to investigate whether this effect is also present in the Salvadoran data, I examined cases of word-initial /s/, of which there were 2943 tokens pertaining to 275 unique words. Of those 275 words, approximately 30% were produced with a lenited allophone at least once. The ten most-frequently lenited lexical items containing a word-initial /s/ appear in Table 63.

Table 63

Most frequently-lenited words beginning with /s/

Word	Total # of occurrences	% lenited
centavos ('cents')	15	53.3%
señor ('sir')	25	48%
se (3rd person pronoun)	530	27.2%
cinco ('five')	44	22.7%
sea ('I/he/she/it is')	50	20%
son ('they/y'all are')	134	18.7%
sé ('I know')	96	16.7%
solo ('only')	103	12.6%
si ('if')	246	9.8%
sí ('yes')	453	8.3%

Similar to Lipski's findings in Honduras, the increased rates of lenition in specific, high-frequency words suggests that some speaker groups may have lexicalized their pronunciations. While this lexicalization would ideally be accounted for in the grammatical model, it is much more difficult to implement than a gradient frequency effect, as specific words would require some kind of "low-confusability" indexing. Future research into this line of inquiry is certainly required.

7.7 Conclusion

Throughout this dissertation, I have aimed to provide a comprehensive account of /s/ lenition in El Salvador via a theoretical model that unites sociophonetic methods with phonological theory. Faced with the challenge of finding a unified theory that accounts for lenition patterns in dialects with pervasive weakening beyond coda position, I have proposed a phonetically-based analysis in which the need to reduce articulatory effort cost while preserving important perceptual distinctions drives variation. I have also argued that, while these language-internal factors establish basic constraints in the grammar, the relative importance of preserving perceptual cues varies for different social groups; this is implemented in the grammar by scaling Faithfulness constraint weights up or down. I have chosen to model the Salvadoran data via maxent, which assigns weights to constraints instead of strictly ranking them. Maxent not only allows for variation but maximizes it, assigning probabilities to all possible output candidates according to shares of total harmony. This maximum entropy approach is particularly suited to modeling Salvadoran /s/ lenition as even the most unlikely outputs—such as deletion in word-initial, intervallic contexts—are observed in the data, albeit very infrequently. Furthermore, like a human language learner, maxent generalizes patterns from the data it observes and is then able to make predictions about novel inputs. While more labor-intensive than traditional statistical analysis, an important benefit of modeling data in maxent is the ability to couch the analysis within phonological theory, as it is able to implement an OT-style analysis for data with free variation.

The findings of this dissertation have been numerous. With respect to Markedness, the grammar shows that the openness of segments flanking /s/ is an important factor in determining effort cost and, subsequently, rates of lenition. This is particularly true for intervocalic contexts: the grammar weights LAZY ("reduce effort cost") constraints highest when both flanking vowels are non-high (weight=8.01), slightly less when one vowel is high (weights=7.03 and 7.25), and least when both vowels are high (weight=5.18). Vowel height also plays an important role in C__V and pause-adjacent contexts, with constraint weights for LAZY increasing when the adjacent vowel is [-high]. These findings are consistent with Kirchner's (2004) effort-based approach to lenition: the biomechanical energy required to move the tongue to the articulatory target and produce tightly constricted [s] increases when adjacent segments are more open.

In preconsonantal contexts, the grammar shows that the drive for place and voicing assimilation—both of which reduce the number of gestures required to produce the CC sequence—

are paramount in determining patterns of weakening. The need for /s/ to agree with the following consonant in place is greater than the need to agree in voicing, as reflected in the high ranking of LAZY(vcl_strid_fric, $\_\_C_{[\text{-coronal, +voice}]}$) (weight=9.50) and LAZY(vcl_strid_fric, $\_\_C_{[\text{-coronal, -voice}]}$) (weight=9.39).[43] The need to minimize the number of glottal gestures is reflected in the increased weights for the constraint LAZY(glottal_abduction) when /s/ occurs before a voiced consonant (weights range from 7.26-8.24) as compared to when it precedes a voiceless consonant (weights range from 4.26-5.17). As expected, LAZY(glottal_abduction) was not sensitive to vowel height for any contexts.

With respect to Faithfulness, the grammar shows that the prosodic position in which /s/ occurs determines the likelihood that lenition will occur. In weaker prosodic positions in which phonetic cues are already less salient and risks of lexical confusability are lower—word-finally and syllable-finally—weights for PRESERVE(strident, voice) are lower. In prosodic positions in which phonetic cues are most important—principally, phrase-initially and word-initially, but also syllable-initially to some extent—weights for PRESERVE(strident, voice) are higher. The tonicity of the syllable in which /s/ occurs also determines constraint weights, although this factor appears unimportant word-medially (in both syllable-initial and syllable-final positions). An important weakness of the model, however, is that it doesn't additionally account for the tonicity of the following syllable when /s/ occurs word-finally before a vowel. An improved model would show that both the tonicity of the syllable in which /s/ originates as well as the syllable to which it resyllabifies are important in this position and context.

Weights for more specific Faithfulness constraints—PRESERVE(strident), PRESERVE(voice), PRESERVE(coronal), and PRESERVE(segmental)—serve to model the behavior of

[43] The weights provided here are for contexts in which the preceding vowel is [-high].

the five lenited variants, which range from [z] to [∅]. The grammar reveals that the input specification [+strident], despite being a defining feature of [s], is not particularly crucial in preserving important perceptual distinctions. Instead, it is the specifications for [+coronal] and [+segmental]—which require that the output still be articulated with the front part of the tongue and that the output be some segment, respectively—that matter most in strong prosodic positions. This reflects the observed hierarchy of allophones in phrase-, word-, and syllable-initial positions: [s], [z], and [s^{θ}] (all [+coronal] and [+segmental]) are preferred, followed by [h] and [ɦ] ([-coronal] but [+segmental]), followed by [∅] ([-coronal] and [-segmental]). While [z], and [s^{θ}]—both slightly less Faithful than [s]—do not occur frequently in the data (*n* = 1037, or 7.2% of total tokens), they serve to mitigate effort cost in salient prosodic positions when necessitated by Markedness constraints. Crucially, the grammar is able to account for the patterns in the data summarized in Table 64, which make a compelling argument for the importance of prosodic strength relations:

Table 64

Allophonic variation according to prosodic hierarchy

Prosodic hierarchy according to the highest domain in which /s/ is initial or final	Allophonic variation					
	[s]	[z]	[s$^{\theta}$]	[h]	[ɦ]	[∅]
Phrase-initial	93.7%	0%	4.0%	1.9%	0%	0%
Word-initial	79.8%	5.5%	6.0%	1.0%	7.0%	1.0%
Syllable-initial	77.9%	6.7%	8.0%	1.0%	5.7%	1.0%
Syllable-final	61.9%	0%	1.2%	11.9%	2.0%	22.7%
Word-final	10.4%	1.4%	1.0%	11.9%	17.2%	58.4%
Phrase-final[44]	50.7%	1.0%	1.0%	4.4%	1.0%	42.5%

Table 64 echoes the patterns espoused by the grammar: as prosodic position weakens, [s] is less probable and [h], [ɦ], and [∅] are increasingly likely to occur. Intermediate allophones [z] and [s$^{\theta}$] occur most frequently in quasi-strong positions (syllable-initial followed by word-initial) but at low rates because their utility for reducing effort cost is not enough to justify the loss of a perceptual cue for some groups of speakers (for other social groups, however, the loss of the perceptual cue is less of a preoccupation and these allophones surface at much higher rates). Finally, while the loss of the [-voice] cue does not have serious implications for the grammar as a whole, this change is particularly bad phrase-initially and syllable-finally, the strongest of respective onset and coda syllabic positions.

With respect to social factors, this dissertation has found that region, urban/rural origin, age, gender, and level of education are all important in determining lenition patterns for Salvadoran speakers. Speakers from San Miguel prioritize Faithfulness the least of the regional groups, resulting in higher rates of lenition overall and particularly in more salient prosodic positions.

[44] As I have discussed in various sections of this dissertation, phrase-final position is unique in that the permissibility to lenite is high because it is a word-final position, yet there is also impetus to produce [s] for phrase-final lengthening, which serves to mark the end of the domain.

However, speakers from Sana Ana present the most interesting case in the end, as they appear to not only be undergoing standardization that outpaces even the capital of San Salvador, but also demonstrate unique gender dynamics that are starkly different from those observed in the other two regions. While men are clearly the more standard speakers in the country as a whole and are likely leading the move toward standardization, women in Santa Ana seem to be leading the change from above. While the case in Santa Ana actually affirms Labovian principles while the other regions do not, it remains an interesting outlier in the data. One possible explanation, proposed earlier in this chapter, is increased dialect contact with neighboring Guatemala, a largely /s/-retaining dialect.

Consistent with accounts of /s/ weakening in other dialects, rural speakers prioritize Faithfulness less than urban speakers, although this difference is smaller than expected. This is likely due to increased movement between rural and urban areas since the start of the Salvadoran Civil war. Interestingly, the urban/rural divide matters significantly more for women than men. I have proposed that this is likely due to increased educational and economic achievement gaps between genders in areas with fewer resources. However, as I discuss at length in Section 7.3.4, the role of gender in Salvadoran /s/ lenition is not clear-cut. It appears that—excepting Santa Ana—men are leading a change from above in El Salvador, which contradicts a wide breadth of theory in both the sociolinguistics and Spanish /s/ weakening literature. I have proposed a number of possible explanations for this situation, including the unique social and economic isolation of Salvadoran women and migration patterns of working-age men. Results also show that women strongly prefer [s^{θ}] while men are much more likely to employ [z] to reduce effort cost in strong prosodic positions; I have proposed that this is due to differences in the characteristics indexed by the two variants, which carry covert prestige for some speakers but not others.

Finally, speakers with lower levels of education were found to prioritize Faithfulness less than speakers with more education. Still, education has been given less attention than other social factors throughout this dissertation because speakers were not balanced and this variable therefore served more as a predictor of other factors (i.e., speakers with less education were more likely to be rural, older, and female) than of lenition itself. Furthermore, the causal relationship between education and lenition is well-established and well-understood, so I believe the more interesting question is *what factors have caused rural, female, and older speakers to be less educated than urban, younger, and male speakers?* rather than, *why do speakers with less education lenite /s/ at higher rates?*

This dissertation has made interdisciplinary contributions. To the field of dialectology, it has contributed valuable information about an understudied dialect of Spanish. This dissertation represents the first comprehensive study on the Spanish of El Salvador which, as evidenced, is rich in nonstandard variation and merits much more attention than it has received. I believe my account of Salvadoran /s/ lenition will also be valuable to scholars interested in the Spanish of Honduras, which by all accounts patterns similarly to that of El Salvador yet also lacks empirical linguistic research. Finally, this study is valuable to dialectology in that it provides a nuanced perspective on Salvadoran [s^{θ}], which I argue has been misunderstood and misrepresented in the literature.

For the field of sociolinguistics, this dissertation has presented an in-depth exploration of the social factors that condition Salvadoran /s/ weakening, including striking patterns that represent a departure from the canonical theory. Furthermore, by interpreting variance between social groups as differences in relative Faithfulness constraint weighting, I have created a tangible model in which language-internal factors determine fundamental patterns in a grammar while social factors determine the relative importance of the constraints that govern those patterns. I believe this is a nuanced and valuable approach to understanding the social variation that exists within a larger grammar, which is one of the fundamental goals of sociolinguistics.

Additionally, this dissertation has contributed to the fields of Spanish phonology and phonological theory. While Spanish /s/ lenition has been studied ad nauseam, many accounts fail to make connections between observed patterns and important aspects of phonological theory which, as I argue in this dissertation, are highly reliant on phonetic principles. This dissertation has taken the phenomenon of joint, systematic onset and coda /s/ weakening—whose nature has been widely debated in the field——and provided an analysis couched within theories of phonetically-based phonology.

7.8 Future directions

While this dissertation represents important progress toward understanding /s/ lenition in El Salvador, there is more work to do. Given the perplexing nature of some of the social factor results, a more immersive ethnographic study is likely needed. Exploring the quantitative relationships between social and linguistic variables is informative to an extent, but a deeper and more nuanced understanding of this speech community is required to fully interpret the results presented here. In this same vein, perceptual studies such as matched-guise tests should be conducted to tease apart the relationship between gender and [s^{θ}]/[z] in order to identify the indexical fields associated with these variants.

Additionally, longitudinal studies are needed to cultivate a better understanding of the /s/ weakening variable in El Salvador. While comparisons of my data with those of Lipski (1984) suggest a shift toward lenition over the past 30 years, a comparison of younger and older speakers' production alleges the opposite. In Section 7.4, I argued that this apparent contradiction is likely the result of comparing apples to oranges: two datasets comprised of very different numbers and populations. For this reason, it is important that future studies of Salvadoran /s/ weakening use comparable samples of speakers and methodologies, and that observed trends are corroborated by longitudinal panel studies that follow the same group of participants over time (Wagner, 2012).

With respect to the variables examined in this dissertation, it is well-known that both local speech rate and register are important conditioning factors for consonant lenition. These factors also fit well into the framework I have proposed: the need to reduce effort cost increases with speech rate for very clear physiological reasons, while the desire to speak clearly and, as an extension, preserve perceptual distinctions increases with register. Both of these factors play a role in Kirchner's (2004) account of consonant lenition (higher speech rates drive up Markedness while higher registers drive up Faithfulness); however, for reasons of feasibility, neither of these variables appears in this dissertation. Future work on Salvadoran /s/ lenition should incorporate these factors as they will undoubtedly improve the model.

Finally, while segmenting and acoustically analyzing tokens for the present project, I observed a fascinating phenomenon in the data that merits further investigation. Assumedly as an overgeneralization of word-final /s/ weakening in prevocalic position, which I have shown to be ubiquitous in El Salvador, some speakers show evidence of sporadic '[h]-intrusion', or epenthesis of the glottal fricative in word-initial position where no underlying /s/ exists, as in *de otros* /deotɾos/ [deˈho.tɾos] ('from others'). This is likely a strategy for hiatus resolution similar to the use of the glottal stop in Nicaraguan Spanish (Chappell, 2013). Given the findings in this dissertation regarding hiatus avoidance word-finally, it will be interesting to investigate the linguistic and social factors that condition Salvadoran [h]-intrusion as an extension of prevocalic /s/ lenition.

Appendix A. Complete list of constraints

Constraint	Constraint type	Definition
LAZY(vcl_strid_fric, V[+high]__V[+high])	Markedness	Assign one violation for every output [s] that occurs between two vowels that are [+high].
LAZY(vcl_strid_fric, V[-high]__V[+high])	Markedness	Assign one violation for every output [s] that follows a vowel that is [-high] and precedes a vowel that is [+high].
LAZY(vcl_strid_fric, V[+high]__V[-high])	Markedness	Assign one violation for every output [s] that follows a vowel that is [+high] and precedes a vowel that is [-high].
LAZY(vcl_strid_fric, V[-high]__V[-high])	Markedness	Assign one violation for every output [s] that occurs between two vowels that are [-high].
LAZY(vcl_strid_fric, V[-high]__C[-coronal, +voice])	Markedness	Assign one violation for every output [s] that follows a vowel that is [-high] and precedes a consonant that is [-coronal] and [+voice].
LAZY(vcl_strid_fric, V[+high]__C[-coronal, +voice])	Markedness	Assign one violation for every output [s] that follows a vowel that is [+high] and precedes a consonant that is [-coronal] and [+voice].
LAZY(vcl_strid_fric, V[-high]__C[+coronal, +voice])	Markedness	Assign one violation for every output [s] that follows a vowel that is [-high] and precedes a consonant that is [+coronal] and [+voice].
LAZY(vcl_strid_fric, V[+high]__C[+coronal, +voice])	Markedness	Assign one violation for every output [s] that follows a vowel that is [+high] and precedes a consonant that is [+coronal] and [+voice].
LAZY(vcl_strid_fric, V[-high]__C[-coronal, -voice])	Markedness	Assign one violation for every output [s] that follows a vowel that is [-high] and precedes a consonant that is [-coronal] and [-voice].
LAZY(vcl_strid_fric, V[+high]__C[-coronal, -voice])	Markedness	Assign one violation for every output [s] that follows a vowel that is [+high] and precedes a consonant that is [-coronal] and [-voice].
LAZY(vcl_strid_fric, V[-high]__C[+coronal, -voice])	Markedness	Assign one violation for every output [s] that follows a vowel that is [-high] and precedes a consonant that is [+coronal] and [-voice].

LAZY(vcl_strid_fric, V[+high]__C[+coronal, -voice])	Markedness	Assign one violation for every output [s] that follows a vowel that is [+high] and precedes a consonant that is [+coronal] and [-voice].
LAZY(vcl_strid_fric, C__V[-high]	Markedness	Assign one violation for every output [s] that follows a consonant and precedes a vowel that is [-high].
LAZY(vcl_strid_fric, C__V[+high])	Markedness	Assign one violation for every output [s] that follows a consonant and precedes a vowel that is [+high].
LAZY(vcl_strid_fric, C__C)	Markedness	Assign one violation for every output [s] that occurs between two consonants.
LAZY(vcl_strid_fric, ##____V[+high])	Markedness	Assign one violation for every output [s] that follows a pause and precedes a vowel that is [+high].
LAZY(vcl_strid_fric, ##____V[-high])	Markedness	Assign one violation for every output [s] that follows a pause and precedes a vowel that is [-high].
LAZY(vcl_strid_fric, V[+high]_____##)	Markedness	Assign one violation for every output [s] that follows a vowel that is [+high] and precedes a pause.
LAZY(vcl_strid_fric, V[-high]_____##)	Markedness	Assign one violation for every output [s] that follows a vowel that is [-high] and precedes a pause.
LAZY(strid_fric, V[+high]__V[+high])	Markedness	Assign one violation for every output [s] or [z] that occurs between two vowels that are [+high].
LAZY(strid_fric, V[-high]__V[+high])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [-high] and precedes a vowel that is [+high].
LAZY(strid_fric, V[+high]__V[-high])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [+high] and precedes a vowel that is [-high].
LAZY(strid_fric, V[-high]__V[-high])	Markedness	Assign one violation for every output [s] or [z] that occurs between two vowels that are [-high].
LAZY(strid_fric, V[-high]__C[-coronal, +voice])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [-high] and precedes a consonant that is [-coronal] and [+voice].

LAZY(strid_fric, V[+high]__C[-coronal, +voice])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [+high] and precedes a consonant that is [-coronal] and [+voice].
LAZY(strid_fric, V[-high]__C[+coronal, +voice])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [-high] and precedes a consonant that is [+coronal] and [+voice].
LAZY(strid_fric, V[+high]__C[+coronal, +voice])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [+high] and precedes a consonant that is [+coronal] and [+voice].
LAZY(strid_fric, V[-high]__C[-coronal, -voice])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [-high] and precedes a consonant that is [-coronal] and [-voice].
LAZY(strid_fric, V[+high]__C[-coronal, -voice])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [+high] and precedes a consonant that is [-coronal] and [-voice].
LAZY(strid_fric, V[-high]__C[+coronal, -voice])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [-high] and precedes a consonant that is [+coronal] and [-voice].
LAZY(strid_fric, V[+high]__C[+coronal, -voice])	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [+high] and precedes a consonant that is [+coronal] and [-voice].
LAZY(strid_fric, C__V[-high]	Markedness	Assign one violation for every output [s] or [z] that follows a consonant and precedes a vowel that is [-high].
LAZY(strid_fric, C__V[+high])	Markedness	Assign one violation for every output [s] or [z] that follows a consonant and precedes a vowel that is [+high].
LAZY(strid_fric, C__C)	Markedness	Assign one violation for every output [s] or [z] that occurs between two consonants.

LAZY(strid_fric, ##____V[+high])	Markedness	Assign one violation for every output [s] or [z] that follows a pause and precedes a vowel that is [+high].
LAZY(strid_fric, ##____V[-high])	Markedness	Assign one violation for every output [s] or [z] that follows a pause and precedes a vowel that is [-high].
LAZY(strid_fric, V[+high]_____##)	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [+high] and precedes a pause.
LAZY(strid_fric, V[-high]_____##)	Markedness	Assign one violation for every output [s] or [z] that follows a vowel that is [-high] and precedes a pause.
LAZY(oral_gesture, V[+high]__V[+high])	Markedness	Assign one violation for every output [s], [z] or [s^{θ}] that occurs between two vowels that are [+high].
LAZY(oral_gesture, V[-high]__V[+high])	Markedness	Assign one violation for every output [s], [z] or [s^{θ}] that follows a vowel that is [-high] and precedes a vowel that is [+high].
LAZY(oral_gesture, V[+high]__V[-high])	Markedness	Assign one violation for every output [s], [z] or [s^{θ}] that follows a vowel that is [+high] and precedes a vowel that is [-high].
LAZY(oral_gesture, V[-high]__V[-high])	Markedness	Assign one violation for every output [s], [z] or [s^{θ}] that occurs between two vowels that are [-high].
LAZY(oral_gesture, V[-high]__C[-coronal, +voice])	Markedness	Assign one violation for every output [s], [z] or [s^{θ}] that follows a vowel that is [-high] and precedes a consonant that is [-coronal] and [+voice].
LAZY(oral_gesture, V[+high]__C[-coronal, +voice])	Markedness	Assign one violation for every output [s], [z] or [s^{θ}] that follows a vowel that is [+high] and precedes a consonant that is [-coronal] and [+voice].
LAZY(oral_gesture, V[-high]__C[+coronal, +voice])	Markedness	Assign one violation for every output [s], [z] or [s^{θ}] that follows a vowel that is [-high] and precedes a consonant that is [+coronal] and [+voice].

LAZY(oral_gesture, V[+high]__C[+coronal, +voice])	Markedness	Assign one violation for every output [s], [z] or [s^θ] that follows a vowel that is [+high] and precedes a consonant that is [+coronal] and [+voice].
LAZY(oral_gesture, V[-high]__C[-coronal, -voice])	Markedness	Assign one violation for every output [s], [z] or [s^θ] that follows a vowel that is [-high] and precedes a consonant that is [-coronal] and [-voice].
LAZY(oral_gesture, V[+high]__C[-coronal, -voice])	Markedness	Assign one violation for every output [s], [z] or [s^θ] that follows a vowel that is [+high] and precedes a consonant that is [-coronal] and [-voice].
LAZY(oral_gesture, V[-high]__C[+coronal, -voice])	Markedness	Assign one violation for every output [s], [z] or [s^θ] that follows a vowel that is [-high] and precedes a consonant that is [+coronal] and [-voice].
LAZY(oral_gesture, V[+high]__C[+coronal, -voice])	Markedness	Assign one violation for every output [s], [z] or [s^θ] that follows a vowel that is [+high] and precedes a consonant that is [+coronal] and [-voice].
LAZY(oral_gesture, C__V[-high]	Markedness	Assign one violation for every output [s], [z] or [s^θ] that follows a consonant and precedes a vowel that is [-high].
LAZY(oral_gesture, C__V[+high])	Markedness	Assign one violation for every output [s], [z] or [s^θ] that follows a consonant and precedes a vowel that is [+high].
LAZY(oral_gesture, C__C)	Markedness	Assign one violation for every output [s], [z] or [s^θ] that occurs between two consonants.
LAZY(oral_gesture, ##____V[+high])	Markedness	Assign one violation for every output [s], [z] or [s^θ] that follows a pause and precedes a vowel that is [+high].

LAZY(oral_gesture, ##____V[-high])	Markedness	Assign one violation for every output [s], [z] or [s$^{\theta}$] that follows a pause and precedes a vowel that is [-high].
LAZY(oral_gesture, V[+high]_____##)	Markedness	Assign one violation for every output [s], [z] or [s$^{\theta}$] that follows a vowel that is [+high] and precedes a pause.
LAZY(oral_gesture, V[-high]_____##)	Markedness	Assign one violation for every output [s], [z] or [s$^{\theta}$] that follows a vowel that is [-high] and precedes a pause.
LAZY(glottal_abduction, V[+high]__V[+high])	Markedness	Assign one violation for every output [s], [h] or [s$^{\theta}$] that occurs between two vowels that are [+high].
LAZY(glottal_abduction, V[-high]__V[+high])	Markedness	Assign one violation for every output [s], [h] or [s$^{\theta}$] that follows a vowel that is [-high] and precedes a vowel that is [+high].
LAZY(glottal_abduction, V[+high]__V[-high])	Markedness	Assign one violation for every output [s], [h] or [s$^{\theta}$] that follows a vowel that is [+high] and precedes a vowel that is [-high].
LAZY(glottal_abduction, V[-high]__V[-high])	Markedness	Assign one violation for every output [s], [h] or [s$^{\theta}$] that occurs between two vowels that are [-high].
LAZY(glottal_abduction, V[-high]__C[-coronal, +voice])	Markedness	Assign one violation for every output [s], [h] or [s$^{\theta}$] that follows a vowel that is [-high] and precedes a consonant that is [-coronal] and [+voice].
LAZY(glottal_abduction, V[+high]__C[-coronal, +voice])	Markedness	Assign one violation for every output [s], [h] or [s$^{\theta}$] that follows a vowel that is [+high] and precedes a consonant that is [-coronal] and [+voice].

LAZY(glottal_abduction, V[-high]__C[+coronal, +voice])	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that follows a vowel that is [-high] and precedes a consonant that is [+coronal] and [+voice].
LAZY(glottal_abduction, V[+high]__C[+coronal, +voice])	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that follows a vowel that is [+high] and precedes a consonant that is [+coronal] and [+voice].
LAZY(glottal_abduction, V[-high]__C[-coronal, -voice])	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that follows a vowel that is [-high] and precedes a consonant that is [-coronal] and [-voice].
LAZY(glottal_abduction, V[+high]__C[-coronal, -voice])	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that follows a vowel that is [+high] and precedes a consonant that is [-coronal] and [-voice].
LAZY(glottal_abduction, V[-high]__C[+coronal, -voice])	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that follows a vowel that is [-high] and precedes a consonant that is [+coronal] and [-voice].
LAZY(glottal_abduction, V[+high]__C[+coronal, -voice])	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that follows a vowel that is [+high] and precedes a consonant that is [+coronal] and [-voice].
LAZY(glottal_abduction, C__V[-high]	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that follows a consonant and precedes a vowel that is [-high].
LAZY(glottal_abduction, C__V[+high])	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that follows a consonant and precedes a vowel that is [+high].
LAZY(glottal_abduction, C__C)	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that occurs between two consonants.
LAZY(glottal_abduction, ##____V[+high])	Markedness	Assign one violation for every output [s], [h] or [s^{θ}] that follows a pause and precedes a vowel that is [+high].

LAZY(glottal_abduction, ##____V[-high])	Markedness	Assign one violation for every output [s], [h] or [s$^{\theta}$] that follows a pause and precedes a vowel that is [-high].
LAZY(glottal_abduction, V[+high]_____##)	Markedness	Assign one violation for every output [s], [h] or [s$^{\theta}$] that follows a vowel that is [+high] and precedes a pause.
LAZY(glottal_abduction, V[-high]_____##)	Markedness	Assign one violation for every output [s], [h] or [s$^{\theta}$] that follows a vowel that is [-high] and precedes a pause.
LAZY(any_gesture)	Markedness	Assign one violation for every output [s], [z], [s$^{\theta}$], [h], or [ɦ] regardless of context.
PRESERVE(strident, voice)/phrase-initial, tonic	Faithfulness	Assign one violation for every non-[s] output in phrase-initial, tonic position.
PRESERVE(strident, voice)/phrase-initial, atonic	Faithfulness	Assign one violation for every non-[s] output in phrase-initial, atonic position.
PRESERVE(strident, voice)/word-initial, tonic	Faithfulness	Assign one violation for every non-[s] output in word-initial, tonic position.
PRESERVE(strident, voice)/word-initial, atonic	Faithfulness	Assign one violation for every non-[s] output in word-initial, atonic position.
PRESERVE(strident, voice)/syllable-initial, tonic	Faithfulness	Assign one violation for every non-[s] output in syllable-initial, tonic position.
PRESERVE(strident, voice)/syllable-initial, atonic	Faithfulness	Assign one violation for every non-[s] output in syllable-initial, atonic position.
PRESERVE(strident, voice)/syllable-final, tonic	Faithfulness	Assign one violation for every non-[s] output in syllable-final, tonic position.
PRESERVE(strident, voice)/syllable-final, atonic	Faithfulness	Assign one violation for every non-[s] output in syllable-final, atonic position.
PRESERVE(strident, voice)/word-final, tonic	Faithfulness	Assign one violation for every non-[s] output in word-final, tonic position.
PRESERVE(strident, voice)/word-final, atonic	Faithfulness	Assign one violation for every non-[s] output in word-final, atonic position.

PRESERVE(strident, voice)/phrase-final, tonic	Faithfulness	Assign one violation for every non-[s] output in phrase-final, tonic position.
PRESERVE(strident, voice)/phrase-final, atonic	Faithfulness	Assign one violation for every non-[s] output in phrase-final, atonic position.
PRESERVE(strident)/phrase-initial, tonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in phrase-initial, tonic position.
PRESERVE(strident)/phrase-initial, atonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in phrase-initial, atonic position.
PRESERVE(strident)/word-initial, tonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in word-initial, tonic position.
PRESERVE(strident)/word-initial, atonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in word-initial, atonic position.
PRESERVE(strident)/syllable-initial, tonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in syllable-initial, tonic position.
PRESERVE(strident)/syllable-initial, atonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in syllable-initial, atonic position.
PRESERVE(strident)/syllable-final, tonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in syllable-final, tonic position.
PRESERVE(strident)/syllable-final, atonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in syllable-final, atonic position.
PRESERVE(strident)/word-final, tonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in word-final, tonic position.
PRESERVE(strident)/word-final, atonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in word-final, atonic position.
PRESERVE(strident)/phrase-final, tonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in phrase-final, tonic position.
PRESERVE(strident)/phrase-final, atonic	Faithfulness	Assign one violation for every output [s^θ], [h], [ɦ], or [Ø] in phrase-final, atonic position.
PRESERVE(voice)/phrase-initial, tonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [Ø]in phrase-initial, tonic position.

PRESERVE(voice)/phrase-initial, atonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in phrase-initial, atonic position.
PRESERVE(voice)/word-initial, tonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in word-initial, tonic position.
PRESERVE(voice)/word-initial, atonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in word-initial, atonic position.
PRESERVE(voice)/syllable-initial, tonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in syllable-initial, tonic position.
PRESERVE(voice)/syllable-initial, atonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in syllable-initial, atonic position.
PRESERVE(voice)/syllable-final, tonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in syllable-final, tonic position.
PRESERVE(voice)/syllable-final, atonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in syllable-final, atonic position.
PRESERVE(voice)/word-final, tonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in word-final, tonic position.
PRESERVE(voice)/word-final, atonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in word-final, atonic position.
PRESERVE(voice)/phrase-final, tonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in phrase-final, tonic position.
PRESERVE(voice)/phrase-final, atonic	Faithfulness	Assign one violation for every output [z], [ɦ], or [∅]in phrase-final, atonic position.
PRESERVE(coronal)/phrase-initial, tonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [∅]in phrase-initial, tonic position.
PRESERVE(coronal)/phrase-initial, atonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [∅]in phrase-initial, atonic position.
PRESERVE(coronal)/word-initial, tonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [∅]in word-initial, tonic position.
PRESERVE(coronal)/word-initial, atonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [∅]in word-initial, atonic position.

PRESERVE(coronal)/syllable-initial, tonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [Ø]in syllable-initial, tonic position.
PRESERVE(coronal)/syllable-initial, atonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [Ø]in syllable-initial, atonic position.
PRESERVE(coronal)/syllable-final, tonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [Ø]in syllable-final, tonic position.
PRESERVE(coronal)/syllable-final, atonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [Ø]in syllable-final, atonic position.
PRESERVE(coronal)/word-final, tonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [Ø]in word-final, tonic position.
PRESERVE(coronal)/word-final, atonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [Ø]in word-final, atonic position.
PRESERVE(coronal)/phrase-final, tonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [Ø]in phrase-final, tonic position.
PRESERVE(coronal)/phrase-final, atonic	Faithfulness	Assign one violation for every output [h], [ɦ], or [Ø]in phrase-final, atonic position.
PRESERVE(segmental)/phrase-initial, tonic	Faithfulness	Assign one violation for every output [Ø] in phrase-initial, tonic position.
PRESERVE(segmental)/phrase-initial, atonic	Faithfulness	Assign one violation for every output [Ø] in phrase-initial, atonic position.
PRESERVE(segmental)/word-initial, tonic	Faithfulness	Assign one violation for every output [Ø] in word-initial, tonic position.
PRESERVE(segmental)/word-initial, atonic	Faithfulness	Assign one violation for every output [Ø] in word-initial, atonic position.
PRESERVE(segmental)/syllable-initial, tonic	Faithfulness	Assign one violation for every output [Ø] in syllable-initial, tonic position.
PRESERVE(segmental)/syllable-initial, atonic	Faithfulness	Assign one violation for every output [Ø] in syllable-initial, atonic position.
PRESERVE(segmental)/syllable-final, tonic	Faithfulness	Assign one violation for every output [Ø] in syllable-final, tonic position.

PRESERVE(segmental)/syllable-final, atonic	Faithfulness	Assign one violation for every output [∅] in syllable-final, atonic position.
PRESERVE(segmental)/word-final, tonic	Faithfulness	Assign one violation for every output [∅] in word-final, tonic position.
PRESERVE(segmental)/word-final, atonic	Faithfulness	Assign one violation for every output [∅] in word-final, atonic position.
PRESERVE(segmental)/phrase-final, tonic	Faithfulness	Assign one violation for every output [∅] in phrase-final, tonic position.
PRESERVE(segmental)/phrase-final, atonic	Faithfulness	Assign one violation for every output [∅] in phrase-final, atonic position.

References

Aguirre, Fausto. 2000. *El Español Del Ecuador*. Loja, Ecuador: Universidad Técnica Particular de Loja.

Alba, Orlando. 1982. "Función Del Acento En El Proceso de Elisión de La /s/ En La República Dominicana." *El Español Del Caribe. Ponencias Del VI Simposio de Dialectología*, 15–26.

———. 1990. *Variación Fonética y Diversidad Social En El Español Dominicano de Santiago.* Santiago, República Dominicana: Pontificia Universidad Católica Madre y Maestra.

Al-Essa, Aziza. 2008. "Najdi Speakers in Hijaz: A Sociolinguistic Investigation of Dialect Contact in Jeddah." Doctoral Dissertation, Colchester, England: University of Essex.

———. 2009. "When Najd Meets Hijaz: Dialect Contact in Jeddah." In *Arabic Dialectology*, edited by Enam Al-Wer and Rudolf de Jong, 203–22. Amsterdam: Brill.

Alfarez, Gabriela. 2000. "Sound Change in a Regional Variety of Cuban Spanish." Doctoral Dissertation, Michigan State University.

Alvar, Manuel. 1959. *El Español Hablado En Tenerife*. Madrid: CSIC.

———. 1995. "Muestras de Polimorfismo En El Español de Argentina." In *La Lengua y Su Expansión En La Época Del Tratado de Tordesillas*, 125–45. Valladolid: Junta de Castilla y León.

Alvar, Manuel, Antonio Llorente, Gregorio Salvador, and José Mondéjar. 1976. *Atlas Lingüístico y Etnográfico de Andalucía (ALEA)*. Vol. VI. Granada.

Al-Wer, Enam. 2014. "Language and Gender in the Middle East and North Africa." In *The Handbook of Language, Gender, and Sexuality*, edited by Susan Ehrlich, Miriam Meyerhoff, and Janet Holmes, 396–411. Hoboken, US: John Wiley & Sons, Inc.

Ávila Muñoz, Antonio Manuel. 1994. "Variación Reticular e Individual de s/z En El Vernáculo Urbano Malagueño: Datos Del Barrio de Capuchinos." *Analecta Malacitana* 17: 343–67.

Azcúnaga López, Raúl Ernesto. 2010. "Fonética Del Español Salvadoreño." In *El Español Hablado En América Central: Nivel Fonético*, edited by Miguel Angel Quesada Pacheco, 83–113. Madrid: Iberoamericana Editorial Vervuert.

Berko, Jean. 1958. "The Child's Learning of English Morphology." *Word* 14 (2–3): 150–77.

Blevins, Juliette. 1995. "The Syllable in Phonological Theory." In *Handbook of Phonological Theory*, edited by John Goldsmith, 204–44. Oxford: Blackwell.

———. 2004. *Evolutionary Phonology: The Emergency of Sound Patterns*. Cambridge: Cambridge University Press.

Blevins, Juliette, and Andrew Garrett. 2004. "The Evolution of Metathesis." In *Phonetically Based Phonology*, 117–56. Cambridge: Cambridge University Press.

Boersma, Paul. 1997. "How We Learn Variation, Optionality, and Probability." In *Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam*, 21:43–58.

Boersma, Paul, and David Weenink. 2018. *Praat: Doing Phonetics by Computer*. http://www.praat.org.

Brogan, Franny, and Mariška Bolyanatz. To appear. "A Sociophonetic Account of Onset /s/ Weakening in Salvadoran Spanish: Instrumental and Segmental Analyses." *Langauge Variation and Change*.

Brown, Earl, and Esther Brown. 2012. "Syllable-Final and Syllable-Initial /s/ Reduction in Cali, Colombia: One Variable or Two?" In *Colombian Varieties of Spanish*, edited by Richard File-Muriel and Rafael Orozco, 89–106. Madrid: Iberoamericana Editorial Vervuert.

Brown, Esther. 2005. “New Mexican Spanish: Insight into the Variable Reduction of ‘La Ehe Inihial’ (/S-/).” *Hispania* 88 (4): 813–24.

Brown, Esther, and Rena Torres Cacoullos. 2002. “Qué Le Vamoh Aher?: Taking the Syllable out of Spanish /s/ Reduction.” In *University of Pennsylvania Working Papers in Linguistics: Papers from NWAV 30*. Vol. 8.

———. 2003. “Spanish /s/: A Different Story from Beginning (Initial) to End (Final).” *Amsterdam Studies in the Theory and History of Linguistic Science Series* 4: 21–38.

Bybee, Joan. 2002. “Word Frequency and Context of Use in the Lexical Diffusion of Phonetically Conditioned Sound Change.” *Language Variation and Change* 14: 261–90.

Calle, Ana María. 2010. “El Fonema /s/, ¿sordo o Sonoro? Un Estudio de Dos Dialectos Ecuatorianos.” *Revista Pucara* 22: 187–206.

Campos-Astorkiza, Rebeka. 2014. “Sibilant Voicing Assimilation in Peninsular Spanish as Gestural Blending.” In *Variation within and across Romance Languages*, edited by Marie-Hélène Côte, Éric Mathieu, and Shana Poplack, 17–38. John Benjamins Press.

Canfield, D. L. 1953. “Andalucismos en la Pronunciación Salvadoreña.” *Hispania 36*, 32–33.

———. 1960. “Observaciones sobre el Español Salvadoreña.” *Filología 6*, 29–76.

———. 1981. *Spanish Pronunciation in the Americas*. Chicago: University of Chicago Press.

Caravedo, Rocío. 1992. “¿Restos de La Distinción /s/ y /θ/ En El Español de Perú?” *Revista de Filología Española* 72: 639–54.

Cardona Ramírez, Mauricio Andrés. 2010a. “Fonética Del Español de Panamá.” In *El Español Hablado En América Central: Nivel Fonético*, edited by Miguel Angel Quesada Pacheco, 177–209. Madrid: Iberoamericana Editorial Vervuert.

———. 2010b. “La Fonética Del Español En Belice.” In *El Español Hablado En América Central: Nivel Fonético*, edited by Miguel Angel Quesada Pacheco, 21–48. Madrid: Iberoamericana Editorial Vervuert.

Carney, P.J., and K.L. Moll. 1971. “A Cinefluorographic Investigation of Fricative Consonant-Vowel Coarticulation.” *Phonetica* 23: 193–202.

Carvalho, Ana María. 2006. “Spanish (s) Aspiration as a Prestige Marker on the Uruguayan-Brazilian Border.” *Spanish in Context 3*, 85–114. https://doi.org/10.1075-sic.3.1.07car.

Casali, Roderic F. 1998. *Resolving Hiatus*. New York: Garland.

Cedergren, Henrietta. 1973. “The Interplay of Social and Linguistic Factors in Panama.” Unpublished doctoral Dissertation, Ithica, NY: Cornell University.

Cepeda, Gladys. 1990. “La Alofonía de /s/ en Valdivia (Chile).” *Estudios Filológicos* 25: 5–16.

———. 1995. “Retention and Deletion of Word-Final /s/ in Valdivian Spanish (Chile).” *Hispanic Linguistics 6/7*, 329–53.

Chappell, Whitney. 2013. “Social and Linguistic Factors Conditioning the Glottal Stop in Nicaraguan Spanish.” Doctoral Dissertation, Columbus, OH: The Ohio State University.

———. 2016. “On the Social Perception of Intervocalic /s/ Voicing in Costa Rican Spanish.” *Language Variation and Change* 28: 357–378.

Chitoran, Ioana, José Ignacio Hualde, and Oana Niculescu. 2015. “Gestural Undershoot and Gestural Intrusion: From Perception Errors to Historical Sound Change.” In *Errors by Humans and Machines in Multimedia, Multimodal, Multilingual Data Processing*. Sinaia, Romania, September 11-13, 2015.

Chomsky, Noam, and Morris Halle. 1968. *The Sound Pattern of English*. New York: Harper & Row.

Cid-Hazard, Susana Margarita. 2003. "Style Variation in Relation to the Phonological Variable /s/ in the Spanish of Santiago, Chile." *Southwest Journal of Linguistics 22*, 13–43.

Coetzee, Andries. 2016. "A Comprehensive Model of Phonological Variation: Grammatical and Non-Grammatical Factors in Variable Nsasal Place Assimilation." *Phonology* 33 (2): 211–46.

Coetzee, Andries, and Shigeto Kawahara. 2013. "Frequency Biases in Phonological Variation." *Natural Langauge and Linguistic Theory 31*, 47–89.

"Corpus Del Español Del Siglo XXI (CORPES)." n.d. Real Academia Española. http://www.rae.es.

Cotton, Elanor, and John Sharp. 1988. *Spanish in the Americas*. Washington D.C.: Georgetown University Press.

Davidson, Justin. 2014. "A Comparison of Fricative Voicing and Lateral Velarization Phenomena in Barcelona: A Variationist Approach to Spanish in Contact with Catalan." In *Romance Languages and Linguistic Theory 2012: Selected Papers from 'Going Romance' Leuven 2012*, edited by K. Lahousse and S. Marzo, 223–244. Philadelphia: Benjamins.

Davies, Mark. 2016. "Corpus Del Español: Two Billion Words, 21 Countries." 2016. http://www.corpusdelespanol.org/web-dial/.

Dohotaru, Puica. 1998. "/S/ Final de Sílaba Y de Palabra En El Habla Juvenil Habanera." *Anuario de Lingüística Hispana* XVI: 149–216.

Eckert, Penelope. 1989. *Jocks and Burnouts: Social Categories and Identity in the High School*. New York: Teach. Coll. Press.

———. 1997. “Age as a Sociolinugistic Variable.” In *Handbook of Sociolinguistics*, edited by Florian Coulmas, 151–67.

———. 2000. *Linguistic Variation and Social Practice*. Oxford: Blackwell.

———. 2012. “Three Waves of Variation Study: The Emergence of Meaning in the Study of Variation.” *Annual Review of Anthropology* 41: 87–100.

“El Salvador National Education Profile.” 2014. Education Policy and Data Center. 2014. https://www.epdc.org/sites/default/files/documents/EPDC%20NEP_El%20Salvador.pdf.

“Encuesta de Hogares de Propósitos Múltiples.” 2007. Census. Ciudad Delgado, El Salvador: Gobierno de la República de El Salvador Ministerio de Economía Dirección General de Estadística y Censos.

Erker, Daniel. 2010. “A Subsegmental Approach to /s/ Weakening in Dominican Spanish.” *International Journal of the Sociology of Language* 2010 (203): 9–26.

———. 2012. “An Acoustically Based Sociolinguistic Analysis of Variable Coda /s/ Production in the Spanish of New York City.” Doctoral Dissertation, New York University.

Face, Timothy L. 2002. “Re-Examining Spanish ‘Resyllabification.’” In *Current Issues in Romance Languages: Selected Papers from the 29th Linguistic Symposium on Romance Languages*, edited by Teresa Satterfield, Christina Tortora, and Diana Cresti, 81–94. Amsterdam: Benjamins.

Fant, Gunnar. 1960. *Acoustic Theory of Speech Production*. The Hague: Mouton & Co.

File-Muriel, Richard. 2009. “The Role of Lexical Frequency in the Weakening of Syllable-Final Lexical /s/ in the Spanish of Barranquilla, Colombia.” *Hispania* 92 (2): 348–60.

File-Muriel, Richard, and Earl Brown. 2010. “The Gradient Nature of S-Lenition in Caleño Spanish.” *University of Pennsylvania Working Papers in Linguistics* 16 (2). http://repository.upenn.edu/pwpl/vol16/iss2/7.

Fontanella de Weinberg, María Beatriz. 1974. “Aspectos Sociolinguisticos del Uso de -s en el Español Bonaerense.” *Orbis* 23: 85–98.

García, Christina. 2015. “Production and Perception of Highland Ecuadorian Spanish.” Doctoral Dissertation, Columbus, OH: The Ohio State University.

García-Amaya, Lorenzo J. 2008. “Variable Norms in the Production of /θ/ in Jerez de La Frontera, Spain.” In *IUWPL7: Gender in Language: Classic Questions, New Contexts*, edited by Jason F. Siegel, Traci C. Nagle, Amandine Lorente-Lapole, and Julie Auguer. Bloomington, IN: IULC Publications.

“Gender Data Portal | Country - El Salvador.” 2018. The World Bank. 2018. http://datatopics.worldbank.org/gender/country/el-salvador.

Gerfen, Chip. 2001. “A Critical View of Licensing by Cue: The Case of Andalusian Spanish.” In *Segmental Phonology in Optimality Theory*, edited by Linda Lombardi, 183–205. Cambridge: Cambridge University Press.

Goldwater, Sharon, and Mark Johnson. 2003. “Learning OT Constraint Rankings Using a Maximum Entropy Model.” In *Proceedings of the Workshop on Variation within Optimality Theory*, edited by Jennifer Spenader, Anders Eriksson, and Dahl, 111–20. Stockholm University.

Gordon, Elizabeth M. 1997. “Sex, Speech and Stereotypes: Why Women Use Prestige Forms More than Men.” *Language in Society* 26: 47–63.

Guitart, Jorge. 1976. *Markedness and a Cuban Dialect of Spanish*. Washington D.C.: Georgetown University Press.

Haeri, Niloofar. 1987. "Male/Female Differences in Speech: An Alternating Interpretation." In *Variation in Language: NWAV-XV at Stanford: Proceedings of the Fifteenth Annual Conference on New Ways of Analyzing Variation,* edited by Keith M. Denning, Sharon Inkelas, John R. Rickford, and Faye McNair-Knox, 173–83. Stanford: Stanford University Press.

Hammond, Robert. 1980. "Las Realizaciones Fonéticas Del Fonema /s/ En El Español Cubano Rápido de Miami." In *Dialectología Hispanoamericana, Estudios Actuales*, edited by Gary E. Scavnicky, 8–15. Washington, D.C.: Georgetown University Press.

Harris, James. 1993. "Integrity of Prosodic Constituents and the Domain of Syllabification Rules in Spanish and Catalan." In *The View from Building 20: Essays in Linguistics in Honor of Sylvain Bromberger*, edited by Kenneth Locke Hale and Samuel Jay Keyser, 177–93. Cambridge, MA: MIT Press.

Hayes, Bruce. 2009. "Manual for Maxent Grammar Tool."

Hayes, Bruce, Robert Kirchner, and Donca Steriade, eds. 2004. *Phonetically Based Phonology*. Cambridge: Cambridge University Press.

Hayes, Bruce, and Donca Steriade. 2004. "Introduction: The Phonetic Bases of Phonological Markedness." In *Phonetically Based Phonology*, edited by Bruce Hayes, Robert Kirchner, and Donca Steriade, 1–33. Cambridge: Cambridge University Press.

Hayes, Bruce, and Colin Wilson. 2008. "A Maximum Entropy Model of Phonotactics and Phonotactic Learning." *Linguistic Inquiry* 39: 379–440.

Hernández Torres, Ramón Augusto. 2010. "Fonética eel Español de Honduras." In *El Español Hablado En América Central: Nivel Fonético*, edited by Miguel Angel Quesada Pacheco, 115–36. Madrid: Iberoamericana Editorial Vervuert.

Hualde, José Ignacio. 1992. "On Spanish Syllabification." In *Current Studies in Spanish Linguistics*, edited by Hector Campos and Fernando Martínez-Gil, 475–93. Washington D.C.: Georgetown University Press.

———. 2005. *The Sounds of Spanish*. Cambridge University Press.

Hualde, José Ignacio, and Monica Prieto. 2002. "On the Diphthong/ Hiatus Contrast in Spanish: Some Experimental Results." *Linguistics* 40: 217–34.

Hundley, James. 1987. "Functional Constraints on Plural Marker Deletion in Peruvian Spanish." *Hispania* 70: 891–94.

Hura, Susan L., Björn J. Lindblom, and Randy Diehl. 1992. "On the Role of Perception in Shaping Phonological Assimilation Rules." *Language and Speech* 35 (1–2): 59–72.

Jäger, Gerhard, and Anette Rosenbach. 2006. "The Winner Takes It All - Almost: Cumulativity in Grammatical Variation." *Linguistics* 44 (5): 937–71.

Johnson, Keith. 2012. *Acoustic and Auditory Phonetics*. 3rd edition. Oxford: Blackwell.

Jongman, Allard, Ratree Wayland, and Serena Wong. 2000. "Acoustic Characteristics of English Fricatives." *Journal of the Acoustical Society of America* 108 (3): 1252–63.

Jun, Jongho. 2004. "Place Assimilation." In *Phonetically Based Phonology*, edited by Bruce Hayes, Robert Kirchner, and Donca Steriade, 58–86. Cambridge: Cambridge University Press.

Kaisse, Ellen. 1999. “Resyllabification Precedes All Segmental Rules: Evidence from Argentinian Spanish.” In *Formal Perspectives on Romance Linguistics*, edited by Marc Authier, Barbara Bullock, and Lisa Reed. Amsterdam: Benjamins.

Keating, Patricia. 2006. “Phonetic Encoding of Prosodic Structure.” In *Speech Production: Models, Phonetic Processes and Techniques*. New York: Psychology Press.

Keating, Patricia, Peggy MacEachern, Aaron Shryock, and Sylvia Dominguez. 1994. “A Manual for Phonetic Transcription: Segmentation and Labeling of Words in Spontaneous Speech.” *UCLA Working Papers in Phonetics* 88: 91–120.

Kingston, John, and Randy Diehl. 1994. “Phonetic Knowledge.” *Language* 70: 419–53.

Kirchner, Robert. 1998. “An Effort-Based Approach to Consonant Lenition.” Doctoral Dissertation, University of California, Los Angeles.

———. 2004. “Consonant Lenition.” In *Phonetically Based Phonology*, edited by Bruce Hayes, Robert Kirchner, and Donca Steriade. Cambridge: Cambridge University Press.

Labov, William. 1966. *The Social Stratification of English in New York City*. Cambridge University Press.

———. 1969. “Contraction, Deletion, and Inherent Variability of the English Copula.” *Language* 45: 715–62.

———. 1972. *Sociolinguistic Patterns*. Philadelphia: University of Pennsylvania Press.

———. 2001. *Principles of Linguistic Change: Social Factors*. Vol. 2. Malden and Oxford: Blackwell.

Ladefoged, Peter. 2003. *Phonetic Data Analysis: An Introduction to Fieldwork and Instrumental Techniques*. Malden, MA: Blackwell.

Lafford, Barbara. 1986. "Valor Diagnóstico Social del Uso de Ciertas Variantes de /s/ en el Español de Cartegena, Colombia." In *Estudios Sobre La Fonología Del Español Del Caribe*, edited by R. Nuñez Cadeño, J. Guitart, and I. Páez Urdaneta, 53–74. Caracas: La Casa de Bello.

———. 1989. "Is Functionalism a Fact? Data from the Caribbean." *Hispanic Linguistics* 3: 49–74.

Lasarte Cervantes, María de la Cruz. 2010. "Datos para la Fundamentación Empírica de la Escisión Fonemática Prestigiosa de /Θs/ en Andalucía." *Nueva Revista de Filología Hispánica* 58 (2): 483–516.

Legendre, G., Y. Miyata, and Paul Smolensky. 1990. "Harmonic Grammar – A Formal Multi-Level Connectionist Theory of Linguistic Well-Formedness: Theoretical Foundations." In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 388–395. Mahwah, NJ: Lawrence Erlbaum Associates.

Lehiste, Ilse. 1972. "Timing of Utterances and Linguistic Boundaries." *Journal of the Acoustical Society of America* 51: 2108–24.

Lenz, Rodolfo. 1940. "Estudios Chilenos (Fonética Del Castellano de Chile)." In *Biblioteca de Dialectología Hispanoamericana*, edited by Amado Alonso. Buenos Aires: Universidad de Buenos Aires.

Lindblom, B. 1983. "Economy of Speech Gestures." In *The Production of Speech*, edited by P. McNeilage, 217–45. New York: Springer Verlag.

Lipski, John. 1984. "On the Weakening of /s/ in Latin American Spanish." *Zeitschrift Fur Dialektologie Und Linguistik* 51: 31–43.

———. 1985. "/S/ in Central American Spanish." *Hispania* 86 (1): 143–49.

———. 1986. “Reduction of Spanish Word-Final /s/ and /n/.” *Canadian Journal of Linguistics 31*, 139–56.

———. 1989. “/S/-Voicing in Ecuadoran Spanish: Patterns and Principles of Consonantal Modification.” *Lingua* 79: 49–71.

———. 1994. *Latin American Spanish*. New York: Longman.

———. 1995. “Blocking of Spanish /s/-Aspiration: The Vocalic Nature of Consonantal Disharmony.” *Hispanic Linguistics 6/7*, 287–327.

———. 1999. “The Many Faces of Spanish /s/-Weakening: (Re)Alignment and Ambisyllabicity.” In *Advances in Hispanic Linguistics*, edited by Javier Gutiérrez-Rexach and Fernando Martínez-Gil. Somerville, MA: Cascadilla Press.

———. 2000. “El Español que se habla en El Salvador y su Importancia para la Dialectología Hispanoamericana.” *Científica 2*, 61–88.

López Morales, Humberto. 1983. *Estratificación Social Del Español de San Juan de Puerto Rico*. Mexico, D.F.: Universidad Nacional Autónoma De México, Instituto de Investigaciones Filológicas.

Lynch, Andrew. 2009. “A Sociolinguistic Analysis of Final /s/ in Miami Cuban Spanish.” *Language Sciences* 31: 767–90.

Mack, Sara. 2010. “Perception and Identity: Stereotypes of Speech and Sexual Orientation in Puerto Rican Spanish.” In *Selected Proceedings of the 12th Hispanic Linguistics Symposium*, 136–47. Somerville, MA: Cascadilla Proceedings Project.

Mason, Keith. 1994. *Comerse las Eses: A Selective Bibliographic Survey of S Aspiration and Deletion in Dialects of Spanish*. Ann Arbor: University of Michigan.

McCarthy, John. 1988. “Feature Geometry and Dependency: A Review.” *Phonetica 43*, 84–108.

McKinnon, Sean. 2012. "Intervocalic /s/ Voicing in Catalonian Spanish." Honors thesis, Columbus, OH: The Ohio State University.

Melguizo Moreno, Elisabeth. 2007. "La Variación de /Θs/: Estudio Comparativo de Dos Muestras de Población de Granada." *ELUA* 21 (1): 1–16.

Méndez Dosuna, Julián. 1985. "La Duración de s en los Grupos SP, ST, SK: A Propósito del Orden Regular de Difusión en Algunos Cambios Fonéticos." In *Symbolae Ludovico Mitxelena Suptuagenario Oblatae*, edited by J.L. Melena, 647–55. Vitoria, Spain: Universidad del País Vasco.

Minnick Fox, Michelle Annette. 2006. "Usage-Based Effects in Latin American Spanish Syllable-Final /s/ Lenition." Doctoral Dissertation, University of Pennsylvania.

Moreno Fernández, Francisco. 2005. "Sobre La Existencia de [θ] En El Español de América." In *Filología y Lingüística. Estudios Ofrecidos a Antonio Quilis*, 1089–1109. Madrid: CSIC-UNED-Universidad de Valladolid.

Moya Corral, Juan Antonio, and Emilio García-Wiedemann. 1995. *El Habla de Granada y sus Barrios*. Granada: Universidad de Granada.

Munson, Benjamin, and Lindsey J. Zimmerman. 2006. "Perceptual Bias and the Myth of the Gay Lisp." presented at the Annual Meeting of the American Speech-Language-Hearing Association, Miami, FL.

Navarro Tomás, Tomás, A.M. Espinosa, and L. Rodríguez-Castellano. 1933. "La Frontera Del Andaluz." *Revista de Filología Española* 20: 225–77.

Ohala, John J. 1981. "The Listener as a Source of Sound Change." In *Papers from the Parasession on Language and Behavior, Chicago Linguistic Society*, edited by Carrie S.

Masek, Roberta A. Hendrick, and Mary Frances Miller, 178–203. Chicago: Chicago Linguistics Society.

———. 1990. "The Phonetics and Phonology of Aspects of Assimilation." In *Papers on Laboratory Phonology I: Between the Grammar and the Physics of Speech*, edited by John Kingston and Mary Beckman, 258–75. Cambridge: Cambridge University Press.

Ohala, John J., and Carol Riordan. 1979. "Passive Vocal Tract Enlargement During Voiced Stops." In *Speech Communication Papers*, edited by J.J. Wolf and D.H. Klatt, 89–92. New York: Acoustical Society of America.

Parodi, Claudia. 2005. "Normatividad y Diglosia En Los Angeles: Un Modelo de Contacto Lingüístico." In *Normatividad y Uso Lingüístico*, edited by Fulvia Colombo and María Angeles Soler, 160–83. México, D.F.: UNAM.

———. 2011. "El Otro México: Español Chicano, Koineización y Diglosia En Los Angeles, California." In , edited by Pedro Martín Butragueño, 217–43. Mexico, D.F.: El Colegio de México.

Penny, Ralph. 2000. *Variation and Change in Spanish*. Cambridge: Cambridge University Press.

Pierrehumbert, Janet. 2001. "Exemplar Dynamics: Word Frequency, Lenition, and Contrast." In *Frequency Effects and the Emergence of Lexical Structure*, edited by J. Bybee and P. Hopper, 137–57. Amsterdam: Benjamins.

Pizarro, Jorge Martínez. 2002. *Los Datos Censales Para Un Análisis Comparativo de La Migración Internacional En Centro América*. Santiago, Chile: The United Nations.

"Población y Estadísticas Demográficas." 2013. Dirección General de Estadística y Censos. 2013. http://www.digestyc.gob.sv/index.php/temas/des/poblacion-y-estadisticas-demograficas.html.

Poplack, Shana. 1979. "Function and Process in a Variable Phonology." Doctoral Dissertation, University of Pennsylvania.

———. 1980. "The Notion of the Plural in Puerto Rican Spanish: Competing Constraints on (s) Deletion." In *Locating Language in Time and Space*, edited by William Labov, 55–67. New York: Academic Press.

Prince, Allen, and Paul Smolensky. 1993. *Optimality Theory*. Oxford: Blackwell.

Quesada Pacheco, Miguel Angel, ed. 2010. *El Español Hablado En América Central: Nivel Fonético*. Madrid: Iberoamericana Editorial Vervuert.

Quesada Pacheco, Miguel Angel, and Luis Vargas Vargas. 2010. "Rasgos Fonéticos Del Español de Costa Rica." In *El Español Hablado En América Central: Nivel Fonético*, edited by Miguel Angel Quesada Pacheco, 155–75. Madrid: Iberoamericana Editorial Vervuert.

Quilis, Antonio. 1993. *Tratado de Fonología y Fonética Españolas*. Madrid: Editorial Gredos.

R Core Team. 2017. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. http://www.R-project.org/.

Ranson, Diana. 1992. "Nominal Number Marking in Andalusian Spanish in the Wake of /s/ Deletion." *Hispanic Linguistics* 4: 301–27.

Regan, Brendan. forthcoming. "A Study of Ceceo Variation in Western Andalusia (Huelva)." *Studies in Hispanic and Lusophone Linguistics*.

———. 2017. "The Effect of Dialect Contact and Social Identity on Fricative Demerger." Doctoral Dissertation, University of Texas at Austin.

Robe, Stanley. 1960. *The Spanish of Rural Panama*. Berkeley and Los Angeles: University of California Press.

Rosales Solís, María Auxiliadora. 2010. “El Español de Nicaragua.” In *El Español Hablado En América Central: Nivel Fonético*, 137–53. Madrid: Iberoamericana Editorial Vervuert.

Ruiz-Sánchez, Carmen. 2004. “El Comportamiento de La /s/ Implosiva En El Habla de Caracas.” *Boletín de Linguística 21*, 48–65.

Sadiqi, Fatima. 2003. *Women, Gender and Language in Morocco*. Leiden: Brill.

Sánchez, Rosaura. 1983. *Chicano Discourse: Socio-Historic Perspectives*. Houston: Arte Público Press.

Schulz, Patrick. 2007. *Schematic Image of Optimality Theory*.

Shadle, Christine H., and Celia Scully. 1995. “An Articulatory-Acoustic-Aerodynamic Analysis of [s] in VCV Sequences.” *Journal of Phonetics* 23: 53–66.

Shewchuk, Jonathan Richard. 1994. “An Introduction to the Conjugate Gradient Method without the Agonizing Pain.” Carnegie Mellon University.

Simons, Gary F., and Charles D. Fennig, eds. 2017. “Summary by Language Size.” *Ethnologue: Languages of the World*. Dallas: SIL International. https://www.ethnologue.com/statistics/size.

Smith, Neilson. 1973. *The Acquisition of Phonology: A Case Study*. Cambridge: Cambridge University Press.

Solé, Maria-Josep. 2003. “Aerodynamic Characteristics of Onset and Coda Fricatives.” In *Proceedings of the 15th International Congress of Phonetic Sciences*, edited by Maria-Josep Solé, Daniel Recasens, and Joaquín Romero, 2761–64. Barcelona: Causal.

“Sounds of Speech.” n.d. University of Iowa Phonetics. Accessed October 24, 2017. http://soundsofspeech.uiowa.edu/index.html#spanish.

Tabain, Marija. 2001. "Variability in Fricative Production and Spectra: Implications for the Hyper- and Hypo- and Quantal Theories of Speech Production." *Language and Speech* 44 (1): 57–94.

Tassara, Gilda, and Carlos Duque. 1987. "Variación Del Segmento /s/ En El Español Culto de Valparaíso." In *VII Seminario de Investigación Y Enseñanza de La Lingüística de La Sociedad Chilena de Lingüística*, 309–324. Valparaíso: Universidad Católica de Valparaíso.

Terrazas, Aaron. 2009. "Middle-Class Immigrant Households in the United States." Migrationpolicy.Org. April 15, 2009. https://www.migrationpolicy.org/article/middle-class-immigrant-households-united-states.

Terrell, Tracy. 1977. "Constraints on the Aspiration and Deletion of Final /s/ In Cuban and Puerto Rican Spanish." *Bilingual Review/La Revista Bilingüe* 4 (1/2): 35–51.

———. 1979. "Final /s/ in Cuban Spanish." *Hispania* 62 (4): 599–612.

———. 1986. "La Desaparición de /s/ Posnuclear a Nivel Léxico En El Habla Dominicana." Edited by R. Nuñez Cadeño, I. Páez Urdaneta, and J. Guitart. *Estudios Sobre La Fonología Del Español Del Caribe*, 117–34.

Terrell, Tracy, and Bernard Tranel. 1978. "Parallelisms between Liason in French and /S/ Aspiration and Deletion in Caribbean Spanish Dialects." *Etudes Linguistiques Sure Les Langues Romanes, Recherches Linguistiques Montreal* 10: 31–50.

"The World Factbook." 2018. Central Intelligence Agency. 2018. https://www.cia.gov/library/publications/the-world-factbook/geos/print_es.html.

Torreblanca, Máximo. 1986. "La 's' Sonora Prevocálica En El Español Moderno." *Thesaurus* 41: 59–69.

Trudgill, Peter. 1972. "Sex, Covert Prestige and Linguistic Change in the Urban British English of Norwich." *Language in Society 1*, 179–95.

Univaso, Pedro, Miguel Martínez Soler, and Jorge A. Gurlekian. 2014. "Variabilidad Intra- e Inter-Hablante de la Fricativa Sibilnte /s/ en el Español de Argentina." *Estudios de Fonética Experimental*, 95–124.

"Urban Population (% of Total) | Data." 2018. The World Bank. 2018. https://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS?end=2016&locations=SV&start=1960&view=chart.

Utgård, Katrine. 2010. "El Español de Guatemala." In *El Español Hablado en América Central: Nivel Fonético*, edited by Miguel Angel Quesada Pacheco, 49–81. Madrid: Iberoamericana Editorial Vervuert.

Valdivieso, Humberto, and Juanita Magaña. 1991. "Variación Fonética de /s/ en el Habla Espontánea." *Revista de Lingüística Teórica y Aplicada* 29: 97–113.

Venables, William N., and Brian D. Ripley. 2002. *Modern Applied Statistics with S* (version Fourth). R. New York: Springer-Verlag.

Villena Ponsoda, Juan Andrés. 1996. "Convergence and Divergence in a Standard-Dialect Continuum: Networks and Individuals in Malaga." *Sociolinguistica* 10: 112–137.

———. 2007. "Interacción de Factores Internos y Externos En La Explicación de La Variación Fonológica. Análisis Multivariante Del Patrón de Pronunciación No Sibilante [θ] de La Consonante Fricativa Coronal /Θs/ En El Español Hablado En Málaga." In *Las Hablas Andaluzas y La Enseñanza de La Lengua. Actas de Las XII Jornadas Sobre La Enseñanza de La Lengua Española*, 69–97. Granada: Universidad de Granada.

Wagner, Suzanne Evans. 2012. “Age Grading in Sociolinguistic Theory.” *Language and Linguistics Compass* 6 (6): 371–82.

Westbury, John, and Patricia Keating. 1986. “On the Naturalness of Stop Consonant Voicing.” *Journal of Linguistics* 22: 145–66.

Wickham, Hadley. 2009. *Ggplot2: Elegant Graphics for Data Analysis*. New York: Springer-Verlag: R.

Yang, Li-Chiung. 2007. “Duration and Pauses as Cues to Discourse Boundaries in Speech.” INTERSPEECH 2007, 8th Annual Conference of the International Speech Communication Association. 458–61. Antwerp, Belgium.

Zimmerer, Frank, and Henning Reetz. 2014. “Do Listeners Recover Deleted Final /t/ in German?” *Frontiers in Psychology*. https://doi.org/doi:10.3389/fpsyg.2014.00735.

Zuraw, Kie. 2009. “Treatments of Weakness in Phonological Theory.” In *Phonological Weakness in English*, edited by Donka Minkova, 9–28. Palgrave MacMillan.