Incomplete Neutralization: A Case of Taiwanese Tone Sandhi

by

Eric Iacono

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in partial fulfillment of the requirements for the degree of

Master of Cognitive Science

in Cognitive Science

Carleton University
Ottawa, Ontario

Abstract

This thesis investigates the perceptual correlates of Taiwanese Tone Sandhi, specifically the extent to which Tone Sandhi causes perceptual neutralization. Two experiments were performed, both of them studies on recognition of checked tone. In the first study speakers of English were asked to distinguish whether pairs of words (which are either the same word, a minimal pair with a difference in tone, or neutralized) heard were the same or different. Accuracy results on this study show that speakers of English correctly categorize minimal pairs and same word pairs at chance, and the neutralized pairs are categorized as same and different in an almost even split. In the second experiment, Taiwanese speakers were given the same task, except the pairs had been modified to all have the same vowel length duration. The unmodified neutralized pairs from the first experiment were also in the test stimuli. Accuracy results show that Taiwanese speakers correctly categorize minimal pairs and same word pairs at a much higher rate compared to English speakers. Taiwanese speakers categorized the unmodified neutralized pairs as being different more often (but not at the same rate as minimal pairs), and the modified pairs as being the same more often (but not as often as same word pairs). This indicates that vowel length is a significant factor in identification of neutralized tones.

Acknowledgements

There are far too many people who were involved in helping me develop this work to its current state to fully list them all. However, I must give immense gratitude to my supervisor Lev Blumenfeld, for his advice and help in making this thesis in it's current readable state, as well as the direction to make sure this research was sensible endeavor. I must also thank my committee, Beth Macleod and Shane Hawkins, for their invaluable comments provided during the defense and for solidifying the core of this thesis.

I would also like to acknowledge Marc Brunelle for planting the initial seed of my research here in Taiwanese tone sandhi, and showing me why tone is a fascinating area of study. I must also thank Grace Kuo, whose research formed a large foundation of the work presented here, and who also provided advice on possible avenues of investigation for this work. The OCP group at University of Ottawa also provided me with excellent discussion and advice for how to proceed on my work.

I would like to thank my good friend Matthew Kelly, for being a constant source of interesting discussion during my time at Carleton, as well as helping to ensure I enjoyed my time there. Lastly, I would like to thank my family, although especially my parents, Sophie and Bob Iacono, as well as my brother and sister-in-law, Mark Iacono and Yi-Wen Wang. Your constant support and generally being a source of comfort when I needed it greatly helped me through this.

Contents

Chapter 1

Introduction

1.0.1 Introduction

This study examines the nature of incomplete neutralization through an analysis of the tone sandhi of Taiwanese, a variety of Southern Min Chinese. Spoken in Southern China, Taiwan, and South-East Asia, Southern Min Chinese languages exhibit five to seven lexical tones (Taiwanese shown in (1)). The q that is seen in the last two examples indicates that the tone exists in a checked syllable (a syllable ending in a voiceless obstruent such as *p*-, *t*-, *k*-, or *ʔ*-).

(1) **Taiwanese Tonal Inventory** (Chen 2000)

Base Form	Pitch Contour	Tone	Gloss
pang	55	High	‘fragrant’
we	24	Rising	‘shoes’
pi	33	Low	‘ailment’
tsu	21	Low Falling	‘house’
hai	53	High Falling	‘ocean’
lip	53q	High Stopped	‘entrance’
bat	31q	Mid Falling	‘know’

In addition, the tones change their quality based on whether they are final or non-final in a Tone Sandhi group domain (a level of prosodic structure). For example, the phrase ‘Many people drive to work’, would be divided into three tonal phrases ‘many people’ /tsin^{55} tse^{33} laŋ^{24}/, ‘drive cars’ /sai^{53} ts^{h}ia^{55}/ and ‘to work’ /k^{h}i^{21} sion^{33} pan^{55}/, and then after sandhi, the sentence would be realized as

(2) Peng (1997)

/tsin^{55} tse^{33} laŋ^{24}/ /sai^{53} ts^{h}ia^{55}/ /k^{h}i^{21} sion^{33} pan^{55}/

↓ ↓ ↓

[tsin$^{\mathbf{33}}$ tse$^{\mathbf{21}}$ laŋ^{24}] [sai$^{\mathbf{55}}$ ts^{h}ia^{55}] [k^{h}i$^{\mathbf{53}}$ sion$^{\mathbf{21}}$ pan^{55}]

The picture becomes more complicated when we consider how such structures can give rise to very ambiguous utterances with identical surface representations. For example, consider the following sentence in Taiwanese which has at least two different meanings:

(3) **Tone sandhi ambiguity** (Kuo 2013)

Surface: *i*^{33} *tiu*^{33} *bin*^{33} *koh*^{53} *khah*^{53} *bai*^{53} *ma*^{21} *m*^{33} *kia*^{55}

Lexical A: “Uncle is not afraid of his face becoming more and more ugly.”

Lexical B: "S/he is not afraid of the scene becoming any worse."

The ambiguity in this example arises from the position of a tonal sandhi group boundary, namely whether one comes after *tiu*^{33} (meaning that it is underlyingly a 33), or whether it comes after *bin*^{33} (in which case, *tiu*^{33} was from an underlying tiu^{24}). Despite this ambiguity, Taiwanese speakers have no trouble understanding words and sentences in Taiwanese. There are two possibilities. One is that tones in domain-final and domain-medial environments, while appearing to be the same contour, are in fact different objects that listeners can distinguish based on some perceptual cues. The other possibility is that there is a perceptual difference associated with the position of the syllable (so if they can recognize the position a syllable is in, they would know what kind of tone it is). This analysis is looked at more in depth in Chapter 2, and forms the basis of the experiments in Chapter 4.

An additional complication here is that how the tones change in Taiwanese present numerous problems for a synchronic analysis. Despite being highly predictable and regular, there does not seem to be a phonological model capable of offering an explanation to account for the patterns observed. This analysis is looked more in depth in Chapter 3.

There are two principal research questions this thesis focuses on. How well do listeners hear the difference between the lexical and sandhi checked (syllables ending in a voiceless obstruent) tones sharing the same surface tone? This is addressed in a perception experiment with English speakers in Experiment 1. The second question is concerning what acoustic cues are Taiwanese listeners using to be able to distinguish between sandhi and lexical tones. Experiment 2 uses a perceptual experiment using modified vowel lengths (the justification is in Section 2.2).

1.1 Structure of the thesis

The thesis will be structured as follows: Chapter 2 will review the incomplete neutralization literature, including the various ways in which the phenomenon has been demonstrated cross-linguistically, as well as the implications of incomplete neutralization on phonology in general; Chapter 3 will discuss the Taiwanese Tone Sandhi phenomenon, and the various ways it has been handled and discussed in the literature; Chapter 4 will discuss two experiments, the first of which will examine English speakers ability to perceive tonal differences in Taiwanese, the second will examine Taiwanese speakers ability to perceive incomplete neutralization after the tones have been digitally altered; finally, Chapter 5 will provide some final conclusions for the thesis.

Chapter 2

Incomplete Neutralization

2.1 Incomplete Neutralization

There is a growing body of work suggesting that some neutralizations in phonology are incomplete. Phonologists in the past had posited that languages were rife with complete neutralization. A classic example is German final devoicing, where an underlying voicing contrast is not preserved, and only voiceless segments surface in coda position. An example of this is shown in (4) (Bloomfield 1933).

(4) **German Coda Devoicing**

/rad/ 'advice' ↘

[rat]

/rat/ 'wheel' ↗

We call this process neutralization, when two underlying distinct segments become phonetically identical with no trace of underlying distinction. Later studies have shown this to be an oversimplification of the process, and that some remnant of the voicing distinction remains on the surface. Port & O'Dell (1985) show that there is an average 15 ms difference in length between the vowel that appears before a 'neutralized' voiceless segment versus one that appears before an underlying vowel in German. Later results (Port & Crawford 1989) showed that there was less aspiration, longer voicing into closure, and marginally shorter closure duration in underlying voiced stops in German, all of which are associated with voiced stops (Lisker 1985). When there exist systematic differences in the neutralized surface forms, we term this incomplete neutralization.

Studies finding incomplete neutralization in final devoicing have been conducted on a number of languages. Most of incomplete neutralization that has been encountered has been only incomplete in production, not perception. That means while there was a systematic difference that was detectable via measurement, this difference was not significant enough to cause speakers to re-categorize the sound heard. In Catalan, native speakers do not perceive a difference, but do produce a difference between underlying voiced and underlying voiceless non-strident obstruents (i.e. *p, t, k*) in word-final position (Dinnsen & Charles-Luce 1987). Slowiaczek & Dinnsen (1985) reported that Polish may be a non-completely neutralizing language, due to some variability in vowel duration in words that ended with an underlying voiceless obstruent and underlying voiced obstruent. However, this durational variability may have been caused by the influence of orthography and ensuing hyperarticulation. This objection is supported by a follow-up perception study which revealed that Polish listeners did not use vowel duration in their voicing judgments (Slowiaczek & Szymanska 1989).

More rarely, speakers are able to perceive a difference in incompletely neutralized pairs of words. A study by van Rooy et al. (2003) showed that Afrikaans speakers are able to produce and perceive a difference between underlying voiced and voiceless final plosives, using a vowel length contrast. Evidence from Warner et al. (2004) in Dutch shows that listeners use durational differences to distinguish words with differing underlying forms, with vowel length being the primary cue. A study by Dmitrieva et al. (2010) in Russian showed that words ending in underlying voiced obstruents have longer vowel duration, shorter closure/frication duration, longer voicing into closure, and shorter release portion, and further studies by Kharlamov (2015) show that listeners can perceive a difference, albeit a very small difference.

A similar phenomenon to that of incomplete neutralization is that of near mergers. These are situations where speakers report that two classes of sounds are the same, but consistently differentiate them in production (Labov et al. 1972). Both incomplete neutralization and near mergers involve an unexpected distinction between either lexical items (in near mergers) or phonemes (in incomplete neutralization) that is not identifiable by traditional phonological methodology. For example, speakers of most forms of British English differentiate words like *source* and *sauce* in production, but report no distinction between them in perception. Similar near mergers have been reported in other varieties of English (e.g. *fool* and *full* in Albuquerque (Di Paolo 1988); *too* vs. *toe* and *beer* vs. *bear* in Norwich (Trudgill 1974); *line* vs. *loin* in Essex (Labov 1971, Nunberg 1980); *meat* vs. *mate* in Belfast (Milroy & Harris 1980, Harris 1985)).

The main puzzle of near merger and incomplete neutralization is how speakers manage to maintain a systematic production difference when they consistently fail to identify the distinction at the conscious or near-conscious level. This is especially perplexing within traditional theories of the phonetics-phonology interface (Cohn 1990, Keating 1990, Pierrehumbert 1990), which assume that phonological representations in the lexicon are categorical, contrastive elements, while the phonetic implementation component computes the degree and timing of articulatory gestures, which are gradient and variable. Information flow is strictly unidirectional in the sense that no articulatory plan can look backward to the phonological encoding, nor can the phonological encoding look back to the lexical level. No lexical information can influence the phonetic implementation directly either, bypassing the level of phonological encoding. On this view, the categorical form of a lexeme wholly determines the phonetic outcome. Phonetic variations on the surface are considered artifacts of the context or performance-induced anomalies. The discovery of systematic subphonemic differences between representations that are otherwise taken to be identical is unexpected.

Responses to the existence of incomplete neutralization and near merger vary widely in the literature. Some phonologists deny the existence of the phenomenon (e.g. Manaster Ramer 1996). Others suggest that the subtle phonetic differences observed are better explained as a consequence of orthographic differences or as variation in speaking style. For example, it has been found that the less the experimental design emphasizes the role of orthography, the smaller the difference in duration (Fourakis & Iverson 1984, Jassem & Richter 1989). Port & Crawford (1989) find that discriminant analysis to classify productions by underlying final voicing is most successful (78% correct) when speakers dictate the words, but least successful (approximately 55% correct) when target words are embedded in sentences that do not draw attention to the minimal pairs (whether read or repeated orally). Conversely, when producing target words in semantically biasing contexts they tend to neutralize more (Charles-Luce 1993, 1997).

While the phonological community reacted to the existence of this phenomenon with caution, the reactions from the sociolinguistic community proved to be much more receptive. Labov (1975), for example, concludes that the existence of near merger suggests that sound change may bring two phonemes into such close approximation that semantic contrasts between them are suspended for native speakers of the dialect, without necessarily leading to merger. In near merger, the underlying category difference has to be supported by something, e.g. contact with another dialect that maintains the distinction (Labov 1994) or, in the case of literate cultures, by orthographic differences (Faber & Di Paolo 1995).

There does exist a middle ground between complete denial and complete acceptance of incomplete neutralization, in which it is restricted to a late stage of phonological production processing, in a module variously termed phonetic implementation or postlexical phonology (Kiparsky 1982). This modular view of the phonology-phonetics interface is traditional, but perhaps the most thorough investigation of its implications for the (incomplete) neutralization issue can be found in Zsiga (1993, 1995, 1997). Zsiga (1995) compared the phonetic behavior of lexical and postlexical palatalization in English; the former is exemplified by *confess*-*confession* (related by semi-productive derivational morphology), the latter by *confess*-*confess your* (a word vs. a compositional phrase with no obvious claim to lexical storage). She found, from both acoustic and electropalatographic data, that the palatals in derived words like *confession* showed no difference from underlying palatals like that in *fresh*, while the palatals in phrases like *confess your* did, since they began like /s/ and only became palatal-like at the end of the segment.

A fundamental problem with the modular approach is the well-known fact that the border between the lexical and postlexical modules has proven to be quite fuzzy. The diagnostics given to distinguish one from the other (Hargus & Kaisse 1993) do co-occur much more often than one would expect to happen by mere accident, but they often disagree. For example, moving German final devoicing out of the lexical phonology to account for its incomplete neutralization is not entirely unproblematic, since this process has been argued to be sensitive to morphological structure (Rubach 1990). Similar points can be made about American English flapping (Steriade 2000) and American English t/d-deletion (Guy 1991).

2.2 Incomplete Neutralization in Tone Sandhi

It has often been claimed that tone in South East Asian Languages is different from tone that occurs in African languages. While African tone languages have small inventories of tones which can be characterized in terms of pitch levels of H (high), M (mid), and L (low), Chinese has a larger number of contrastive tones, a fact recognized in the numbers 1 (low) through 5 (high) widely used in Chinese linguistics to transcribe tones. Second, Chinese tonal inventories contain more tones with complex contours than their African counterparts (Zhiming 2011). When tones occur in a compound or phrase, they may undergo change, called tone sandhi. The term lexical tone is used to refer to tones that are lexically specified for each morpheme and sandhi tone is ued to refer to tones derived from lexical tones in specific contexts. Many of the tone sandhi processes that are common in African languages, such as spreading, are also attested in Chinese (such as the Wu dialect of Shanghaiese). Naturally, tone sandhi processes which involve a contour are either unique to or more common in Chinese and other Asian languages with large tonal inventories. The complexity and range of tone sandhi in Chinese have contributed to our understanding of tonal phenomena in general. Consider the tonal inventory of Beijing Mandarin. Beijing has four lexical tones shown in (5).

(5) **Beijing Mandarin Tonal Inventory** (Chen 2000)

I	55
II	35
III	214
IV	51

The tones are labeled I–IV in keeping with convention of Chinese linguistics labeling tones after their Middle Chinese (ca. 900 CE) categories: *ping* 'even', *shang* 'rising', *qu* 'departing', and *ru* 'entering', respectively. Beijing Mandarin also has tone sandhi as shown in (6).

(6) **Beijing Mandarin Tone Sandhi**

1. 214 → 35 / ___ 214

When there are two adjacent (and adjacency applies over prosodic word boundaries) 214 tones, the first in the pair changes to a 35. This change can be thought of a composite of two processes where the contour changes from concave to rising and the pitch height is raised at the same time. According to the thorough examination of phonetic studies of this pattern given in Myers & Tsay (2003), it seems that contrary to Chao (1948), sandhi rising (35) tone is systematically lower in pitch than lexical rising tone in the speech of native speakers of Beijing Mandarin, consistent with the hypothesis that it preserves aspects of its lexical form (Kratochvil 1968, Zee 1980, Xu 1997). Peng (2000) demonstrates via production and perception experiments that while the 35 that is derived from an underlying 214 is distinct from an underlying 35 tone, native speakers of Mandarin do not perceive this distinction. Peng concludes that Mandarin speakers must hear, learn, and reproduce a fine phonetic distinction between underlying 35 and sandhi 35, that in adulthood they are no longer able to perceive. Thus, Chinese tone sandhi has features of incomplete neutralization.

Yu (2007) discusses a different sort of subphonemic tonal phenomenon in Cantonese, this time involving the near merger of distinct lexical categories rather than incomplete neutralization due to morphophonemic alternation (e.g. tone sandhi). While Cantonese does not possess tone sandhi as defined above, it does possess a similar phenomenon referred to in the literature as *pinjam*, or 'changed tone'. In this, certain morphemes are derived from their respective semantic correlates just by a change of tone. For clarity, the tonal inventory of Cantonese is shown in (7).

(7) **Cantonese Tonal Inventory** (Yu 2007)

si	55	'poetry'
si	33	'to try'
si	22	'affairs'
si	23	'market'
si	35	'to cause'
si	21	'time'
sɪk	55	'to know'
sɛk	33	'to kiss'
sɪk	22	'to eat'

The only clear case of *pinjam* is the derivation of mid-rising tones from semantically related syllables with a non-high-level, non-midrising tone. For example, the nominalization of verbs carrying a non-rising tone is indicated purely by a change in tone (e.g. [sou^{33}] 'to sweep', related to [sou^{35}] 'a broom'). The productivity of this pattern is unclear and many native speakers do not recognize any relationship between the derived forms and their derived base. However, Yu's instrumental analysis demonstrated that the morphologically derived items have an F0, on average, 4.8 Hz higher than the former, though without a difference in contour shape. Moreover, as is typical for near mergers, native speakers listening to isolated syllables cannot reliably distinguish the two categories. Although this case does not shed light directly on tone sandhi, it does seem to demonstrate that it is indeed possible for tonal systems to involve subphonemic lexical contrasts.

The Chinese language that is the focus of this dissertation is Taiwanese, a Chinese language of the South Min language family. Taiwanese has 7 tones, as shown in (8).

(8) **Taiwanese Tonal Inventory** (Chen 2000)

Base Form	**Pitch Contour**	**Tone**	**Gloss**
pang	55	High	fragrant
we	24	Rising	shoes
pi	33	Low	ailment
tsu	21	Low Falling	house
hai	53	High Falling	ocean
lip	53q	High Stopped	entrance
bat	31q	Mid Falling	know

Taiwanese also has a very robust tone sandhi process as seen in (9). Every tone in Taiwanese undergoes this process if it is not at the end of a tone sandhi phrase domain. A tone sandhi phrase domain is not confined to a phonological word or a phonological phrase; its size could vary from a disyllabic phrase to a full sentence with multiple syntactic phrases (Kuo 2013). Also of note is that the tone sandhi process is structure preserving (Kiparsky 1985); every tone that exists in sandhi position is a tone that also exists in lexical position.

(9) **Taiwanese Tonal Inventory** (Chen 2000)

	Base Form		**Sandhi Form**	
a.	*tsin p'ang*	"very fragrant"	*p'ang tsui*	"perfume"
	55		33	
b.	*p'e we*	"leather shoes"	*we tua*	"shoe laces"
	24		21	
c.	*wi pĩ*	"stomach ailment"	*pĩ lang*	"sick person"
	33		21	
d.	*k'i ts'u*	"build a house"	*t'su ting*	"roof"
	21		53	
e.	*tua hai*	"big ocean"	*hai kĩ*	"ocean front"
	53		55	
f.	*ts'ut lip*	"exit and entrance"	*lip k'ao*	"entry point"
	53q		31q	
g.	*m bat*	"not know"	*bat li*	"literate"
	31q		53q	

In the past, the literature has categorized the tone sandhi of Souther Min languages to be categorical and completely neutralizing. Hwei–Bing (1988) does address the question of neutralization in the production of Taiwan Southern Min tone sandhi, basing the analysis on the production and perception of nonsense disyllables. However, speakers were mixed from different dialects (one of the six speakers had a different tone sandhi system from the rest) and the researchers did not systematically control the segmental environment; the use of nonsense syllables rather than real words may also have given rise to somewhat artificial productions.

As a brief aside, it is worthwhile to investigate how Southern Min influences Mandarin in Taiwan. As mentioned above in Peng's (2000) experiments, lexical 35 and sandhi 35 are distinct in Beijing Mandarin. However, studies with speakers of Taiwan Mandarin (the variety of Mandarin spoken in Taiwan and strongly influenced by Taiwanese) generally do not find differences between sandhi 35 and lexical 35, whether in production (Chang & Su 1994, Fon 1997) or in perception (Chang & Su 1994, Peng 2000). Peng (2000) did find that the mean fundamental frequency (F0) of sandhi rising tone was significantly lower than that of lexical fall rise tone, but by a much smaller amount than has been reported for Beijing Mandarin (2.3 Hz, as compared with the 17.5 Hz found by Zee (1980)). Myers & Tsay (2003) suggest that this apparent dialect difference, where Mandarin tone sandhi in Taiwan is categorical while Mandarin tone sandhi in Beijing is not, provides indirect support for the claim that tone sandhi in Taiwan Southern Min is truly categorical, assuming that this property is carried over to the Mandarin acquired by Southern Min speakers as a second language.

An experiment by Myers & Tsay (2008) sought to test the hypothesis that tone sandhi in Taiwanese is completely neutralizing. This was done by having speakers read sentences and measuring several phonetic properties of either 55 or 33, when it appeared in lexical or sandhi position. A consistent difference arose in duration effects. Lexical 33 was on average 61 ms longer than sandhi 33. Since Taiwanese tone sandhi occurs everywhere except at boundaries, then it is possible speakers are associating phrase final lengthening with different tonal categories. Final lengthening is a major temporal cue that listeners use to locate prosodic boundaries, with greater effects of lengthening at successively higher levels of prosodic domains. The final rime in a prosodic domain is lengthened, which reflects a decrease of the articulation rate at the end of the phrase (Beckman & Edwards 1990) . Lengthening is a robust boundary marker in a variety of languages (English: Price et al. 1991, Wightman et al. 1992, Turk & Shattuck-Hufnagel 2007; Mandarin: Shen 1992; Korean: Cho and Keating 2001).

However, there was no significant difference in F0 when produced in lexical or sandhi position. The difference in overall F0 between sandhi 33 and lexical 33 was not only not significant, but also extremely small, only 2.3 Hz. This is quite a bit smaller than the 17.5 Hz difference found by Zee (1980) between lexical 35 and sandhi 35 in Beijing Mandarin, or even than the 4.8 Hz difference Yu (2007) reports for near merger in Cantonese tones. Myers and Tsay did find one case of incomplete neutralization of two tones, which had nothing to do with the prosodic position. One of the outcomes of the Taiwanese tone sandhi process, is an underlying rising tone (24) and an underlying high level tone (55) have the same sandhi tone outcome – a mid level tone (33). Myers and Tsay compared sandhi tone 33 derived from 55 and from 24. This is the only case where two underlyingly different tones might be neutralized into the same surface tone in sandhi context. Their finding was that speakers produced overall higher F0 for sandhi tone derived from 55 than for sandhi tones derived from 24, and the tones derived from 55 were 8 ms longer on average than the tones derived from 24. This result is interesting because the prosodic position of the two sandhi tones is the same, and thus any differences between them must be due to incomplete neutralization rather than phrasal prosody.

It thus seems that there are at least two potential kinds of information that Taiwanese listeners might use in coping with tone sandhi. First, sandhi tones might be phonetically distinct from lexical tones, e.g. in pitch. If a small phonetic difference is perceptible by listeners, such physical differences could be used to recover the lexical identity of a phonetic tone, and there would be no mystery at all in how Taiwanese listeners recognize words in the face of tone sandhi.

The second kind of information would be less direct, concerning the detection of prosodic boundaries. Since sandhi tones occur within the tone sandhi group domain and citation tones occur before all kinds of phrasal boundaries, then information about phrasal boundaries is information about surface tone forms. For example, the tone before a full pause must be a citation form. Once a listener hears a pause, then it becomes clear that, say, a phonetic mid level tone right before the pause should be a lexical tone 33, and not the sandhi form of a lexical 55 or of a lexical 24. This information does not become clear until after the entire syllable is heard. In other words, if this were the only available boundary information, then recognition would be rather delayed. However, other prosodic information as well as contextual information might provide cues to upcoming phrasal boundaries. If a boundary can be anticipated, then tonal recognition need not be delayed.

In her 2013 dissertation, Kuo performed several studies analyzing Taiwanese tone sandhi, specifically trying to test how Taiwanese speakers disambiguate sentences of the type seen in (3). In the first study, participants were presented with increasingly longer segments of a Taiwanese word and were asked to propose (either a forced-choice or an open-choice response) the word being presented, and to give a confidence rating after each fragment. The words used in this experiment were all carrying checked tones, and were extracted from either sandhi or non-sandhi (lexical) position. Kuo claimed that if the neutralization was complete, one would expect the correct identification rate to be equal between sandhi and lexical tones. The larger the disparity in the correct identification rate between sandhi and lexical tones, the more it appears that the neutralization is incomplete. The results of the experiment found correct identification of lexical tones as being 77% and sandhi tones as 65%, indicating a small but significant effect of incomplete neutralization. In addition, there is a significant difference in the time required to identify the tones, with tones in sandhi position taking much shorter time (Kuo conjectures the reason for this stems from sandhi forms being the more common token, and so listeners identified them faster). These findings suggest that some phonetic cues must have coarticulated with the stimuli, and that these cues are relevant to the location of the boundary. The only thing that matters within a tone sandhi group is the cues concerning the location of the boundaries. In other words, lexical and sandhi tones can be neutralized, but only partially, because the occurrence of these tones depends on their position relative to the boundary.

In the second study by Kuo, she examined the acoustic differences between lexical and sandhi tones using a Taiwanese speech corpus. Three potential differences were examined based on previous literature: final lengthening, F0, and creaky voice phonation. Final lengthening (a.k.a. pre-boundary lengthening) refers to the often-observed pattern that the final syllable preceding a prosodic boundary is longer than that syllable in a no-boundary position. F0 magnitude is another salient cue that has been proposed by a number of studies, especially studies with tone languages. Creaky voice phonation often accompanies a low F0 value and could be used as a boundary signal if it occurs more frequently before boundaries. In a study of Mandarin Chinese connected speech, Belotel-Grenié & Grenié (2004) found that the second syllable of a disyllabic word (therefore, at a prosodic boundary) is where the creaky voice was found because the latter syllable is associated with a lower F0 value.

In the study, disyllabic phrases (the corpus was constructed via native speakers reading stories) were examined, with specific attention being paid to contrast tones in lexical versus sandhi position, along the dimensions outlined above. In terms of duration, tones at the end of a prosodic boundary were found to be longer than tones in sandhi position for all the tones by a significant margin. Compared to F0, which only found a significant difference in F0 magnitude for a few tones: 31 has a wider range in lexical position, whereas 33 has a wider range in sandhi position. Kuo concludes that F0 range is not as strong a cue as duration, as not as many significant effects were found. The last dimension is creaky voice, which was found to have some significant results. Tones 51, 31q, 53q, and 33 were found to be significantly creakier in lexical position compared to sandhi position.

While not all the tones differed in length, F0 range, and creaky voice, the fact that some did suggests at least that part of the tonal pattern is incompletely neutralized. Experiment 2 in Chapter 4 will focus on the perception of tones after they have been manipulated such that the vowel length in sandhi and lexical tones in checked syllables are the same. Although Kuo did identify that checked syllables also exhibited a significant difference in creaky voice, vowel length was found to be significant in all tones. Moreover, manipulation of tokens to equalize vowel length to create natural sounding tokens is much easier than equalizing creaky voice.

The following chapter will focus on different theoretical models, namely derivational phonology and Optimality Theory (Prince & Smolensky 2004), and how they have examined the tone sandhi in Taiwanese.

Chapter 3

Taiwanese Tone Sandhi

3.1 Introduction

Taiwanese, like Mandarin as a member of the Sino-Tibetan language family, is a tonal, monosyllabic language. Traditionally, a Taiwanese syllable is decomposed into three parts, an initial, a final, and the tone. For example, the syllable dan^{53} "to wait", has /d/ as an initial, /an/ as a final, and /53/ represents a high falling tone. There are 18 initials, 47 finals, and 7 lexical tones. An initial is always just a consonant, but a final can be further decomposed in one to three vowels and a consonant. And then for each final, there exists a checked variant, where it ends in a /p/, /t/, /k/, or /?/.

The data in (10) shows all the tones that appear in Taiwanese. To read (10), treat the starting number (53 for example) as the lexical tone and then follow the arrows once to find what that tone surfaces as in sandhi position (e.g. 55, if the lexical tone was 53). A 'q' is appended to the structure on the right of the circle to show that these are tones in checked syllables (that is, CVC where the coda is a voiceless stop).

(10) Taiwanese Tone Circle (Chen 2000)

$$
\begin{array}{ccccccccc}
 & & 24 & \rightarrow & 21 & & & & \\
 & \nearrow & & & & \searrow & & & \\
33 & & & & & & 53 & 31q \leftrightarrow 53q & \\
 & \nwarrow & & & & \swarrow & & & \\
 & & & 55 & & & & &
\end{array}
$$

Southern Min tone sandhi is a phrase-level rather than word-internal phenomenon, as has been made well-known by work on its interaction with syntactic structure (Chen 1987, Lin 1994). The pattern involves alternations between tones as they appear in juncture position (i.e. the right edge of a prosodic constituent called a tone group) and in context position (elsewhere). Taiwanese speech can be organized into a hierarchy of three major prosodic levels: Intonation Phrase ("IP", hereafter), Tone Sandhi Group ("TSG", hereafter) and Word/Syllable (Kuo 2013). TSGs are primarily defined syntactically, aligning at the right edge with the right edge of syntactic phrases, unless the phrases are in particular syntactic relations with other elements (namely if the phrase is lexically governed, according to Lin (1994)).

The tone alternations themselves are quite complex, affecting each of the seven contrastive tones differently. Particularly notable (and rather rare even among Sinitic tone sandhi systems, as the survey in Chen (2000) reveals) is the fact that the alternations do not create nonlexical tones (i.e. they are structure-preserving, in the terminology of Kiparsky (1982)).

This chapter examines the various theoretical approaches that have been used to examine Taiwanese Tone Sandhi, starting with series of derivational rules up until variations on the Optimality Theory paradigm. Taiwanese tone sandhi is an unusual phonological process, as the tones form two separate chain shifts that loop around and form chain circles, and have not been found to exist in languages outside of this small branch of the Sino-Tibetan language family.

3.2 Generative Phonology

Starting at the very beginning of the generative framework in The Sound Pattern of English (Chomsky & Halle 1968), we can try to describe this problem using a set of counterfeeding ordered rules. However, we immediately encounter an ordering paradox as you can see below (TS is used to denote the boundary of a Tone Sandhi domain as created by the rule in (11)):

(11) (1)24 → 21 / ___ not TS

(2)33 → 21 / ___ not TS

(3)55 → 33 / ___ not TS

(4)53 → 55 / ___ not TS

(5)21 → 53 / ___ not TS

Based on the above ordering, the prediction is that 21 will never surface in lexical position, which is contradictory to the data above. This is independent of the ordering; no matter the choice, the last input in the rule ordering will never surface. One possible fix here is to posit an intermediate representation for the first rule (e.g. 22 → 21′) and then have a rule at the end that turns the intermediate representation back to its original form (e.g. 21′ → 21). While usage of intermediate representations is not inherently without merit, creation of an intermediary tone that never surfaces when no other phonological process in the language (or any other Chinese dialect) appears to do so (Chen 2000), would make this solution ad-hoc and not well motivated externally.

If features are used, then we can achieve a more succinct solution. Recall from Section 2.2:

(12) **Taiwanese Tonal Inventory**

Pitch Contour	Tone
55	High
24	Rising
33	Low
21	Low Falling
53	High Falling
53q	High Stopped
31q	Mid Falling

Wang's (1967) more elegant solution is using an 'exchange rule' (the values of the two features 'exchange'), shown in (13):

(13)

$$\begin{bmatrix} \alpha & \textsc{high} \\ \beta & \textsc{falling} \end{bmatrix} \rightarrow \begin{bmatrix} \beta & \textsc{high} \\ --\alpha & \textsc{falling} \end{bmatrix} / \ \_\_ \text{ not TS}$$

With one rule and two features, the entirety of Mainstream Taiwanese is captured by this rule. The apparent downside comes in when we consider other dialects of Southern Min that are very close. Consider Longxi, another Southern Min dialect located geographically very close:

(14) Longxi Tone Sandhi (Chen 2000)

313 → 33
↗ ↘
14 11
↖ ↙
41

This looks like an identical tone sandhi process operates in both Taiwanese and Longxi. It would still be easy to dismiss that these two processes share the same structure, since languages share similar looking rules that have no relation. However, if we translate the tones into their historical equivalents (labelled I–III, with a or b used to describe a historical register split), both tone circles collapse into the same structure.

(15) Proto-South Min Tone Sandhi (Chen 2000)

Ib → IIIb
↗ ↘
Ia IIIa
↖ ↙
II

The tonal categories hold constant across dialects, but are filled by phonetic values that vary from one dialect to the other. The probability that the two dialects have converged by chance in having developed independently of each other is extremely unlikely. This strongly implies that the tone sandhi process developed prior to the two dialects diverging phonetically. It would seem that the most fruitful approach would be to study the proto-language that these dialects arose from. Very little is known about this language, and as such we must focus our efforts on what can be determined synchronically.

Derviational phonology does not accurately capture the tone circle, and the next section will move onto attempts at describing the tone circle in the more modern framework of Optimality Theory.

3.3 Optimality Theory

Optimality Theory using only markedness constraints (structural well-formedness) and faithfulness constraints (contrast preserving), hereafter referred to as Classic OT, has been used to describe the tonal pattern of Southern Min languages. Yip (2002) provides the a solution in Optimality Theory to Southern Min tonal shifts, and she does so by positing the following constraints:

(16) Yip OT Constraints (Yip 2002)

1. *Contour: Contour tones are more marked than level tones.
2. *H≫*L: High tones are more marked than low tones.
3. *Rise≫*Fall: Rising tones are more marked than falling tones.
4. Space–100%: Level tones must be 100 per cent of the pitch range apart. (Can only be satisfied by systems with two or fewer evenly spaced tones: [H, L], or any single tone.)
5. Space–50%: Level tones must be 50 per cent of the pitch range apart. (Can only be satisfied by systems with 3 or fewer evenly spaced tones: [H, M, L], [H, M], [H, L], [M, L], or any single tone.)
6. 6T/4T/2T: Distinguish 6/4/2 tonal distinctions

The general idea is that in different positions the inventory is defined by different OT mini-grammars. In each position, different numbers of contrasts are required, which means formally that one of the constraints 6T, 4T or 2T is ranked above tonal markedness constraints. As a shorthand, she collapses 1-3 in (16) into the constraint MinArticEffort (defined as *Contour ≫ *H≫*L ≫ *Rise≫*Fall), and posits the rankings as shown in (17).

(17) Yip OT Rankings

- Lexical position: Space–50% ≫ 2T ≫ 3T ≫ 4T ≫ 5T ≫ MinArticEffort ≫Space–100% ≫ 6T
- Sandhi position: Space–50% ≫ 2T ≫ 3T ≫ 4T MinArticEffort ≫Space–100% ≫ 5T ≫ 6T

In sandhi position, there exist only four tones (21,53,55,33) in unchecked syllables, compared to the five that exist in lexical (24 is the unique tone). Once the inventory is defined for a given context, it constrains the particular output by requiring that the tone be part of the inventory set. It does not, however, explain why a given lexical item chooses the tone that it does.

While the notion of multiple rankings for a single grammar exists and has been proposed elsewhere with good motivation (Kiparsky 2000, Inkelas & Zoll 2007), typically the framework has involved sending a single form through multiple strata (e.g. having 4 different constraint rankings, and sending it through each one before it surfaces), or having multiple cophonologies that have different rankings depending on the morphological environment (e.g. a word would have one ranking for an adverbial form, and another for progressive). However, no such formalism is present in this solution. More importantly, Yip admits that, "the pairings (between tone and syllable) are indeed lexical, a position I have resisted for years but at present see no alternative to," indicating that this model is still missing the ability to generate the correct output given a certain input. It is only able to generate the correct inventory for a given position.

Rather than discuss inventory, Moreton (2004) shows OT to be completely unable to handle circular chain shifts. That is, OT cannot select the correct surface form given any input form when the grammar under analysis has a circle shift. The formal proof will not be replicated here, but the thrust of the argument is the notion of harmonic ascent, i.e. unfaithful candidates only win against faithful candidates if they reduce markedness. If the relation 'more harmonic than' is thought to be transitive (and there does not seem to exist any good reason to think it is not) then we can immediately see the issue when we claim that 21 is more harmonic than 22, which is more harmonic than 44, which is more harmonic than 53 which is more harmonic than 21, which is now a contradiction as we have claimed 21 is more harmonic than itself (or less marked). Moreton argues that Taiwanese is not a counter-example to Harmonic Ascent, and that "if a language has an alternation that cannot be computed by these constraints, the learner can detect this by the failure of the learning algorithm to converge; the deficiency would then be suppleted by adding an undominated ad hoc constraint which simply insisted that a particular lexical item, in a particular environment, must surface in a particular way." This does resolve the issue, it also makes the broad claim that every such circle is a so-called 'crazy rule'. It specifies that a large number of undominated ad hoc constraints must be learned, one for every possible lexical item that can not be learned in OT. This is very similar to the allomorph selection hypothesis (discussed in the next section), in terms of it being a very powerful and unrestrained theory. Much like those theories though, it fails to capture the systematic generalization that, in sandhi context, syllables bearing one tone in the input share the same tone in the output, which always differs from their shared tone in the input.

However, the formal proof of OT's inability to handle circle shifts only applies to Classic OT, and the next section will look at attempts at extending the Classic OT framework to see if expansions can better capture the tone circle and at what cost.

3.4 Alternate Explanations

Many possible expansions to the above theoretical frameworks have been proposed in order to account for this problem. Since circles and chains are definitionally impossible to handle in classical OT, we can either try to expand the power of our theories massively in order to handles circles, or expand it a little in order to handle chains, and then try to reduce the Taiwanese tone circle to a chain. First, let us examine what possible expansions will let us solve chains in OT. One possible avenue that can handle chain shifts is trans-derivational anti-faithfulness (Alderete 2001). Mortensen (2002) makes an attempt at describing other languages (Jingpho and A-Hmao, a Tibeto-Burman and Western Hmongic language, respectively) with tonal chain shifts (although not circles) using anti-faithfulness constraints. Specifically,

- ¬Ident–Contour (which says the feature [contour] must have a different specification in output than in input)
- ¬Ident–Tone (which says this constraint is violated when a tone in the output

shares all of the same tone features as the corresponding tone in the input and is preceded in the output by a [+contour] tone).

A full explanation of how these constraints are used is omitted since it would require a full explanation of the the tonal systems of Jinhpho and A-Hmao (neither of which are related to Southern Min), as well as their historical backgrounds. The most important part is that the solution given by Mortensen provides a good empirical motivation for the existence of anti-faithfulness constraints. There are still two glaring issues with anti-faithfulness. The first is that there is still no principled way to limit the power of what can be described by anti-faithfulness. The second, is that the notion of anti-faithfulness contravenes the markedness/faithfulness interaction that lies at the heart of OT.

Another expansion to OT could be adopting local constraint conjunction of faithfulness constraints (LCC) (Kirchner 1996). In this, two constraints are conjoined and the conjoined constraint is considered to be in violation if and only if both are in violation. An example of how this is constructed is in (18):

(18) Local Constraint Conjunction (Kirchner 1995)

-		Input	*x	F1 & F2	*y	F1	F2
a.	/x/	[x]	*!				
b.	☞	[y]			*	*	
c.		[z]		*!		*	*
a.	/y/	[x]	*!			*	
b.		[y]			*!		
c.	☞	[z]					*

The constraints and inputs are abstract in (18); *x and *y are markedness constraints against allowing x or y to surface, respectively. F1 and F2 are two faithfulness constraints for some arbitrary features. The constraint F1 & F2 is a constraint that is violated if and only if F1 and F2 are violated. [x] and [y] share the same value on F2 and differ on F1; [x] and [z] differ in value on both F1 and F2; [y] and [z] differ on F2 value but share the same F1 value. The tableau in (18) shows two different inputs, /x/ and /y/ being evaluated according to this ranking. The first three rows show [y] being chosen as the winner for /x/, since [z] differs in F1 and F2 value from /x/ F1 & F2, thus violating the conjoined constraint. For the /y/ input, [z] is chosen as the winner, as the markedness constraints dominate the faithfulness constraints here, preventing anything other than [z]. This accurately describes the process of a chain shift in OT terms.

Moreton & Smolensky (2002) provide an argument for incorporating constraint conjunction into OT, well motivated by data from different languages, and show it is possible to capture chain shifts (but not chain circles) non-trivially once LCC is allowed.

Although local constraint conjunction is supported by data, it also encounters some of the same foundational issues as anti-faithfulness above, as well as a computational issue. The current landscape of constraints in OT (Ashley et al. 2010) estimates the current number of universal constraints to be around 1700. Of these, 900 are markedness constraints and 500 are faithfulness constraints. Even if we assume, as Moreton & Smolensky do, that constraints can only conjoin with constraints in similar domains (and that conjoined constraints can not conjoin with other constraints), the introduction of constraint conjunction increases the total number of constraints several orders of magnitude at a conservative estimate.

Even with the above caveats, we can look at ways people have tried to reduce Taiwanese to chains. Zhang et al. (2009) claim that the Taiwanese tone circle is actually illusory, and that what appears is fossilized allomorphy. Justification for this is based on a 'wug-test' experiment testing the productivity of the tone sandhi process. The experiment involved creating a disyllabic utterance out of two different words, with either of the morphemes being an actual morpheme of Taiwanese, an accidental gap (i.e. could exist in the Taiwanese syllabary but does not due to historical reasons), or any combination of the two. Results showed that Taiwanese speakers regularly apply tone sandhi on actual Taiwanese words, but mostly fail to do so on accidental gap phrases. It fails to account for why certain classes of loanwords participate in the process. If the pattern was simply fossilized allomorphy resulting from sound changes that occurred in the histories of these languages, there is no good reason to expect that it should apply to any loanwords at all. Even assuming productivity of a process indicates whether it is present in a grammar, these experiments have been challenged on grounds that the constructions given to participants are not given proper context and explanation. Later experiments show that when Taiwanese speakers are given proper context, tone sandhi becomes a productive process (Chuang et al. 2011).

If allomorphy selection is not sufficient to explain the situation, incomplete neutralization could be. In Section 2.2, we referred to the fact that the tones in Taiwanese could be incompletely neutralized, and what is perceived to be a circle is really due to over-categorization (e.g. the 24 and 33 that surface as 21 in sandhi position could be different on some perceptual level that is salient to a native speakers). If it is the case that speakers are able to make such distinctions (even if it is only between one pair along the circle), then at worst we have a long tone chain, which is resolvable by one of the above OT extensions. We will revisit this at the end of Section 4.3, once we have determined the perceptual properties Taiwanese speakers use (if any) to make this distinction

Chapter 4

Perception Experiments

4.1 Motivation

This chapter presents two experimental studies, which together explore how lexical tones and sandhi tones differ perceptually. In the first study, English listeners were asked to identify whether two tones heard were the same or different, with the aim of measuring the ability of English speakers to identify tone. In the second study, a similar paradigm is used, but all but one of the tone sandhi pairs of Taiwanese were compared acoustically to test for neutralizations.

Kuo (2013), through the use of a gating experiment, demonstrates that Taiwanese speakers can determine whether the word they are listening to was a word in sandhi position or in lexical position, despite having the same tone contour. The experiment also determined that the rate of response for sandhi tones was faster than lexical tones. This indicates that unlike Mandarin speakers, Taiwanese speakers are able to distinguish two surface similar tones. Kuo identifies three variables, based on a corpus study, duration, F0 range, and H1*-H2* (the difference between the amplitude of the first and second harmonic, which is strongly correlated with creaky voice phonation (Belotel-Grenié 2004)) as the main acoustic differences between lexical and sandhi tones. However, of the three variables, only duration is systematically different in sandhi tones and lexical tones, in that lexical tones are found to be longer. As a result, Experiment 2 will focus on the ability of Taiwanese speakers to perceive a difference between lexical and sandhi tones, after the vowel lengths of lexical and sandhi tones have been made identical.

4.2 Experiment 1: English speaker perception of tone sandhi

Experiment 1 tests the ability of English speakers to distinguish surface-identical tones that are lexical vs. those in sandhi form. Recall the inventory of Taiwanese tones and how they interacted in Figure (10) in Chapter 3. If English speakers can make tonal distinctions (especially with regards to making a distinction between lexical and sandhi tones), then we can also test by what cues they are making tonal distinctions, meaning Taiwanese speakers would not be needed (who are relatively rare in Eastern Canada).

This experiment examines checked tones in Taiwanese, comparing for example, a 53 that's an underlying /53/ vs. 53 that's an underlying /31/. Previous studies of Taiwanese tones looked at non-checked tones. The expectation is that English speakers should have difficulty performing well on the task due to lack of exposure to tonal languages and tone sandhi, suggesting that English speakers are not able to make tonal distinctions.

4.2.1 Stimuli

A trained female Taiwanese speaker[1] from the southern part of Taiwan provided the stimuli. Eighteen different checked monosyllabic words (half with high falling /53/, the other with low falling /31/) were extracted from two positions, the first syllable of a disyllabic phrase (sandhi) and the second syllable of a disyllabic phrase (lexical). The other syllable in each disyllabic phrase was controlled as carrying a mid level tone. Then each of these extracted syllables was paired with another extracted syllable that either was identical (in segmental content), identical but with a different tone, or identical but the result of one of the syllables undergoing a tone sandhi process. In total, there were 108 tested pairs: 18 Syllables x 3 types of pairs x 2 orderings. The following table shows the possible pairs:

(19)

		Lexical		Sandhi	
		/53/ → [53]	/31/ → [31]	/53/ → [31]	/31/ → [53]
Lexical	/53/ → [53]	Same	Different	Test Point	X
	/31/ → [31]	Different	Same	X	Test Point
Sandhi	/53/ → [31]	Test Point	X	Same	Different
	/31/ → [53]	X	Test Point	Different	Same

Tones with // are used to indicate lexical tones, and [] are used to indicate sandhi tones. 'Same' indicates that these tones are the same and in the same environment, 'Different' are minimal pairs in the same environment with different tones, and 'Test Point' are tones which are the same contour but in different positions ('X' has been used to denote different tones in different positions, which is not tested in this exper-

[1] Thanks to Grace Kuo for providing the stimuli for both experiments.

iment). The accuracy of 'Test Point' trials are measured as well as the accuracy of 'Same' and 'Different' trials.

4.2.2 Participants

Twenty four native English speakers (13 females, 11 males) from Ottawa, Ontario, age 18-30, with no knowledge of tonal languages were found via social media (Facebook).

4.2.3 Procedure

At the start of the session, each participant receives verbal instructions that they are going to listen to pairs of words in Taiwanese and that they will need to categorize whether they think the tone in these words are the same or different.

In each trial, they would click on a symbol to hear one of the stimuli pairs, and they would be presented with a forced choice between same or different. When they had selected one of the two choices, the next pair in the experiment would be played.

4.2.4 Results

The results are given in 20. The average rate at which listeners correctly identified two different tones as different was 49.0%, and the rate at which they correctly identified two identical tones as same was 50.1%. As for test pairs, where the same surface tone was either the result of sandhi, subjects identified them as the same tone 54% of the time and as different tones 46%. A chi-square test of independence was performed to examine the relation between stimuli type and response. The relation between these variables was not significant, $\chi^2(1,\ 864)= 2.23, p=0.14$ between pairs that had identical tones and pairs that had the same surface tone as either the result of sandhi or as the realized lexical form. A similar non-significant result existed $\chi^2(1, 864)= 1.23, p=0.27$ between pairs that had different tones and pairs that had the same surface tone as either the result of sandhi or as the realized lexical form. This indicates that the experiment did not detect any difference in the way English speakers differentiate tones.

(20) **Experiment 1: Total number of responses, separated by pair type and then by response given.**

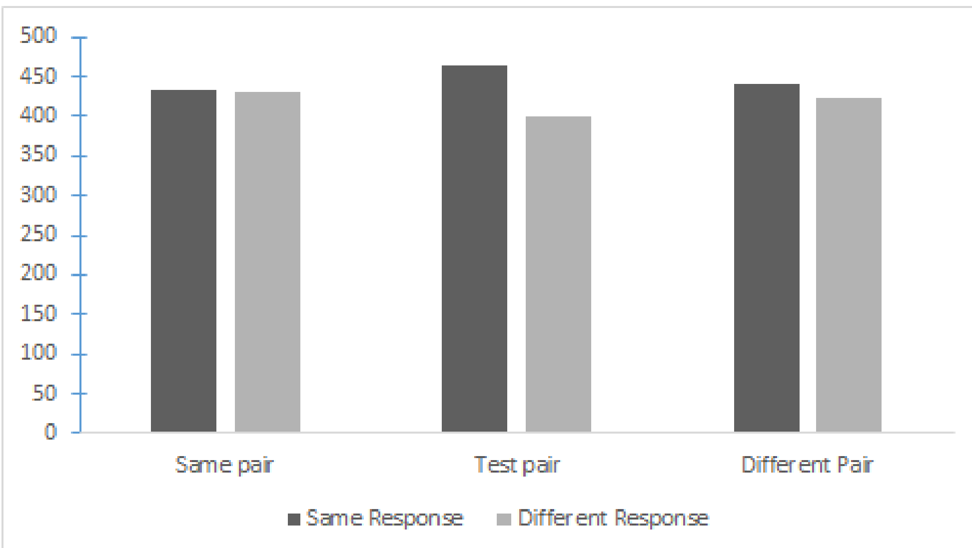

4.2.5 Discussion

The absence of evidence from Experiment 1 suggests that English speakers are unable to distinguish between sandhi and lexical tones, as the accuracy rate is close to chance, and there is no statistically significant difference between pair type on response. English speakers are unable to distinguish incompletely neutralized pairs from same and different pairs and seem to be completely deaf to tone, in contrast to

Taiwanese speakers who have demonstrated the ability to do so (Kuo 2013). This is still unusual, as we would expect that deafness to tone would mean that they never hear a difference and that it should be categorized entirely as being the same. The forced choice nature of the experiment could have lead speakers to believing that they need to categorize some of the tokens as different, and as such, would just randomly categorize the tokens as either.

4.3 Experiment 2: Tone Perception by Taiwanese Listeners

Experiment 2 is designed to examine the specific acoustic cues that Taiwanese speakers use to distinguish sandhi tones from lexical tones using the criteria established by Kuo (2013) in Section 4.1. This experiment takes a set of auditory stimuli in Taiwanese, and using Praat software, manipulates the signal to ensure that the vowel duration is identical between sandhi and lexical tones.

4.3.1 Stimuli

The same as in Experiment 1 except for the following: the length of every tone in a sandhi position (average of 256 ms) and lexical position (average of 681 ms) were measured. Then, the average of these two values were taken (469 ms), and using a praat script, all tones were changed to be this length. As an additional check, the naturalness of all modified tokens was judged by a naive native Taiwanese[2] speakers who ensured that after manipulation, the tokens were still understandable

[2] Thanks to Yi-Wen Wang for being a judge

Taiwanese words. The unmodified neutralized pairs from Experiment 1 were added to the experimental design as well. In total, there were 144 tested pairs, tested twice: 18 Syllables x 4 types of pairs x 2 orderings, making a total of 288 trials.

4.3.2 Participants

Twenty four native Taiwanese speakers (15 females, 9 males), age 18-50, were recruited through Facebook Social Groups and local associations.

4.3.3 Procedure

At the start of the session, each participant receives verbal instructions that they are going to listen to pairs of words in Taiwanese and that they will need to categorize them if they think they are written with the same or different characters.

In each trial, they would click on a symbol to hear one of the stimuli pairs, and they would be presented with a forced choice between same or different. When they have selected one of the two choices, the next pair in the experiment would be played.

4.3.4 Results

The results are given in (21). The average rate at which listeners correctly identified two different tones as different was 72.3%, and the rate at which they correctly identified two identical tones as same was 87.3%. Tone pairs where the tone contours were identical due to tone sandhi and then modified in length to be identical were identified as being same pairs 69.6% and the pairs that were the same but had not been modified to be identical in length were identified as same pairs 44.4% of the time. A chi-square test of independence was performed to examine the relation between pair type and response type. The relation between these variables was significant, between same and modified test pairs $\chi^2(1, 1728) = 160.1$, $p<0.001$, between unmodified test pairs and modified test pairs $\chi^2(1, 1728) = 223.4$, $p<0.001$, and between unmodified test pairs and modified test pairs $\chi^2(1, 1728) = 104.3$, $p<0.001$. An ANOVA analysis was done on the reaction times between the four pair types (means are shown in (22)), with finding a significant results $F(3, 20) = 127.2$, **p**<0.001, η^2 =0.12 indicating a small effect size. A follow up Tukey HSD test was run to evaluate differences between each pair, and significant (p<0.001) differences were found for each possible pair, except for different vs. test pairs, which was a non-significant result (p=0.373).

(21) **Experiment 2: Total number of responses, separated by pair type and then by response given.**

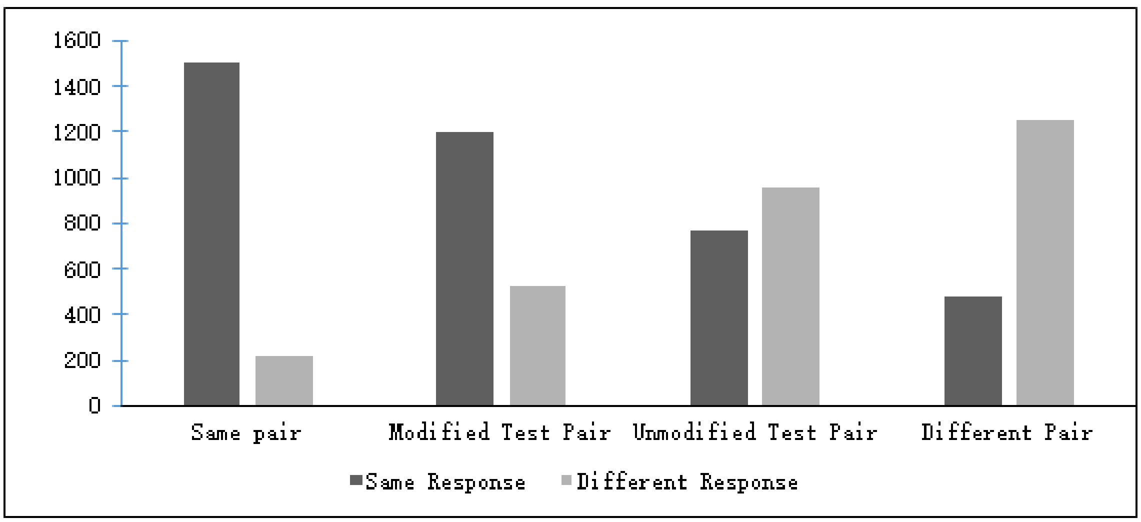

(22) **Experiment 2: Average reaction time in seconds.**

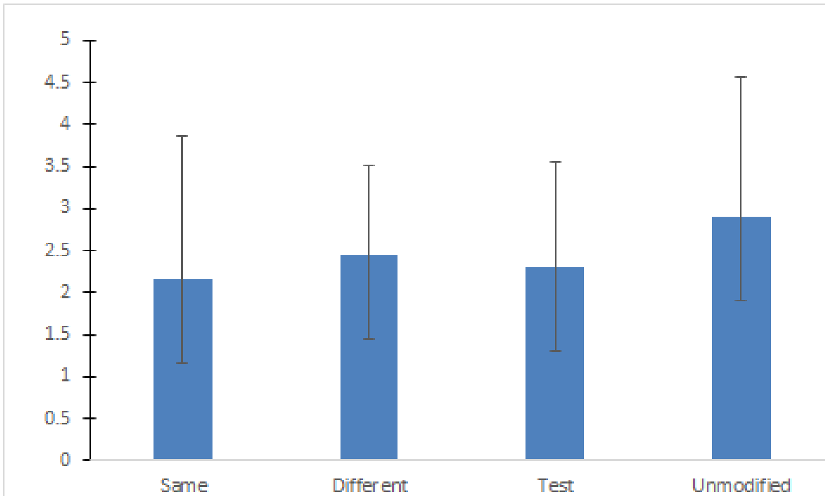

4.3.5 Discussion

This study used the same perception design as in Experiment 1, but returned drastically different results. The Taiwanese speakers much more accurate than English speakers at the same and different pairs. In addition, Taiwanese participants performed differently depending on whether they were presented with modified test or unmodified test pairs. Recall that the only difference between the modified and unmodified test pairs (i.e. pairs which have an identical tone contour due to one of the syllables undergoing tone sandhi) is that length of the sandhi tones was increased in the modified pairs and the vowel lengths of the lexical tones was decreased to a mid-point that was uniform across all the pairs. This was done because, as pointed out in Sections 2.2 and 4.1, Kuo's corpus study identified duration as the only phonetic factor that was consistently different for sandhi and lexical tones. Duration is a crucial cue to identifying the underlying tone. We can conclude this based on the results of Experiment 2, which show a statistically significant difference in performance between modified and unmodifed pairs.

However, there is a significant difference in performance on pairs where the tones are identical and pairs where the tones are identical by Taiwanese speakers, due to one of the tones undergoing tone sandhi changes. This indicates that while vowel length does have a significant effect on the ability of Taiwanese speakers to perceive neutralization, it is not quite the whole difference. Additional factors outlined in Sections 2.2 and 4.1 such as manipulation of F0 and phonation would need to be tested in future work.

The reaction time results fall in line with what has been shown elsewhere in the literature (Li 2015). Reaction time is highly correlated with ambiguity in the acoustic signal, and ambiguity in the signal increases the further a given sound is from an already present category. For a listener of a tonal language, tones that are underlying the same contour and surface as the same contour are categorized easily as being the same, and so speakers categorize them faster. Similarly, tones that are underlying a different contour and surface as different contours are categorized easily as being different, and so speakers categorize them faster. Tones that are underlying a different contour but surface as the same contour have properties of their underlying form still present, and so they are further away from a known category, thus taking longer time to categorize. When some of the properties of the underlying form have been taken away (in the form of vowel length reduction in Experiment 2), they are closer to being a more often heard category, and so are categorized more quickly. The small effect size here indicates that while the effect was significant, it was quite small, and as a result we can conclude that reaction time was only affected by the type of pair the participant was presented with by a small amount.

Chapter 5

Discussion & Conclusion

This thesis examined and tested the process of Taiwanese tone sandhi, by examining its perception and acoustic correlates. The purpose was to answer the questions of “how well do listeners hear the difference between the lexical and sandhi checked tones sharing the same surface tone?”, and “do Taiwanese listeners use vowel duration to distinguish between sandhi and lexical tones?”.

5.1 Perceiving the difference

Taiwanese sandhi tones and citation tones are widely assumed to share the same tonal categories. In other words, pairs of tones are said to have the same tonal output, and thus their pitch contours should show no difference. However, the prosodic positions where the citation tone and the sandhi tone appear are different. Lexical tones appear in the final position of a Tone Sandhi Group whereas sandhi tones are only found in the non-final positions of a Tone Sandhi Group.

The results from the perception experiment in Section 4.2 indicate that English speakers were not even able to identify a difference between two syllables that were minimal pairs, or whether two syllables were identical (with accuracy at chance rate for both). The original prediction was that English speakers would be unable to make the distinctions on this test due to their lack of exposure to tones and tone sandhi, which was borne out by the data.

The Taiwanese speakers in Section 4.3, were able to correctly identify when two presented syllables were the same or different with a much higher accuracy compared to the English speakers. Further research should be done to determine whether the ability to make this distinction was due to the Taiwanese speakers' past exposure to Taiwanese, or whether it is a more general learned ability of being able to make tonal distinctions. For example, a contrast with Mandarin (fewer tonal distinctions, but with tone sandhi present) or Cantonese (more tonal distinctions, but no tone sandhi) speakers on the same test.

5.2 Altered cues

Across languages, domain-final positions usually involve final lengthening, pitch declination and creaky voice. This suggests that there might be differences in these acoustic measures between lexical and sandhi tones. The experiment in Section 4.3 tests part of this hypothesis (specifically, final lengthening as justified in Section 2.2) by having participants attempt to distinguish between pairs of Taiwanese words that have been modified to be equal in length.

While Kuo had demonstrated that Taiwanese speakers distinguish between sandhi and lexical tones, and that there was a measurable difference in production between citation and sandhi tones, there was no data to show that Taiwanese speakers were using the factors she identified. In Experiment 2, Taiwanese speakers were more likely to categorize tone pairs where one of the tones had been neutralized to the other (e.g. a 53q that is underlying 31q being paired with a 53q that is underlying 53q) as being the same when the vowel lengths were modified to be identical. The same types of pairs when there was no modified vowel length were more likely to be categorized as being different tones, in contrast. We can also conclude that the act of modifying the signal itself had little to do with the categorization, given the accuracy of pairs that were the same (e.g. a 53q that is underlying 31q being paired with a 53q that is underlying 31q) or different (e.g. a 53q that is underlying 53q being paired with a 31q that is underlying 31q) were both fairly high and they were modified as well.

While Kuo did demonstrate that Taiwanese speakers perceived sandhi and lexical tones differently, and that there were measurable acoustic differences between sandhi and lexical tones, it was not made clear if those acoustic differences were being used by speakers to distinguish between sandhi and lexical tones. The research done in this thesis demonstrates those acoustic differences are being used by Taiwanese speakers as a means of differentiating between lexical and sandhi tones. Moreover, this research also shows that speakers are using cues of the underlying form in the surface form to lexically identify certain forms. The implication of this is that the tone circle is illusory; tones are not neutralizing categorically, which means that the features of the underlying forms can be coded in the rules or constraints. Recall from Section 3.2 that derivational rules did not work for tone circles, since whatever rule was last would prevent one of the forms from ever surfacing:

(23) **Naive Checked Rule**

(1)53q → 31q / ___ not TS

(2)31q → 53q / __ not TS

This could now be restated with the knowledge that the surface form that the sandhi 'neutralizes' to is not identical to the form that is underlying the same as that form (i.e. 53q that is 53q underlying is not the same as 53q that is underlying 31q):

(24) **Re–stated Checked Rule**

(1)53q → $31q^1$ / __ not TS

(2)31q → $53q^1$ / __ not TS

While this research did demonstrate that vowel length is a cue that is being used to differentiate between lexical and sandhi tones, we can not conclude that vowel length is the sole cue being used. The results indicate there are still factors that cause speakers to be able to categorize incompletely neutralized pairs as being the same at a different rate from pairs where the tone contour and position are the same. Moreover, this data does not provide insight as to whether vowel length is a phonetic cue particular to the tone, or whether vowel length is used by Taiwanese speakers to signify a tone sandhi boundary, and can then use the knowledge of the boundary to reconstruct the underlying form from the surface one. Future research would test this by analyzing Taiwanese speakers' perception of boundaries, and what cues they use to distinguish different types of prosodic boundaries.

Appendices

Appendix A

Stimuli for Experiments 1 & 2

1. Same:
    (a) bak^{31}(L)—bak^{31}(L)
    (b) bak^{31}(S)—bak^{31}(S)
    (c) bak^{53}(L)—bak^{53}(L)
    (d) bak^{53}(S)—bak^{53}(S)
2. Different:
    (a) bak^{53}(S)—bak^{31}(S)
    (b) bak^{53}(L)—bak^{31}(L)
    (c) bak^{31}(S)—bak^{53}(S)
    (d) bak^{31}(L)—bak^{53}(L)
3. Test Points:
    (a) bak^{53}(L)—bak^{53}(S)

(b) *bak*^{31}(L)—*bak*^{31}(S)

(c) *bak*^{53}(S)—*bak*^{53}(L)

(d) *bak*^{31}(S)—*bak*^{31}(L)

Appendix B

Consent Form

See following page.

Carleton
UNIVERSITY
Canada's Capital University

CUREB clearance #:

Consent Form

Title: Perception of Taiwanese Sandhi Tones

Funding Source: Lev Blumenfeld

Date of ethics clearance: To be determined by CUREB (as indicated on the clearance form)

Ethics Clearance for the Collection of Data Expires: To be determined by CUREB (as indicated on the clearance form)

I ____, choose to participate in a study on Perception of Taiwanese Tones. This study aims to study English speakers' perception of tone in other languages. **The researcher for this study is Eric Iacono in the Institute of Cognitive Science.**
He is working under the supervision of Lev Blumenfeld in the Institute of Cognitive Science.

This study involves one 30-60 minute test. With your consent, data about your performance on this test will be collected. You will be made to listen to various pairs of sounds in another language, and you will have to judge whether these sounds are the same or different.

You have the right to end your participation in the study at any time during the session, for any reason. This will result in the data collected from that session being destroyed. At the end of the session, you will be prompted whether you wish to keep your results. If you do not want to keep your results, this will result in the data collected from that session being destroyed. Since all data collection is anonymous, once the test is complete and you have left the room you can not withdraw.

Should you register for a class at a later date with either researcher in a Teacher or TA position, you may request an alternate marker.

As a token of appreciation, you will receive a $5 Tim Horton's gift card. This is yours to keep, even if you withdraw from the study.

All research data and any notes will be kept anonymous. Research data will only be accessible by the researcher and the research supervisor.

This document has been printed on both sides of a single sheet of paper.
Please retain a copy of this document for your records.

Once the project is completed, all research data will be kept for ten years and potentially used for other research projects on this same topic. At the end of ten years, all research data will be securely destroyed. (Electronic data will be erased and hard copies will be shredded.)

If you would like a copy of the finished research project, you are invited to contact the researcher to request an electronic copy which will be provided to you.

The ethics protocol for this project was reviewed by the Carleton University Research Ethics Board, which provided clearance to carry out the research. Should you have questions or concerns related to your involvement in this research, please contact:

CUREB contact information:
Professor Louise Heslop, Chair
Professor Andy Adler, Vice-Chair
Carleton University Research Ethics Board
Carleton University
511 Tory
1125 Colonel By Drive
Ottawa, ON K1S 5B6
Tel: 613-520-2517
ethics@carleton.ca

Researcher contact information:
Eric Iacono
Institute of Cognitive Science
Carleton University
Tel:
Email: ericiacono@cmail.carleton.ca

Supervisor contact information:
Lev Blumenfeld
Institute of Cognitive Science
Carleton University
Tel:
Email: levblumenfeld@cunet.carleton.ca

____________________________ ________________
Signature of participant Date

___________________________ ________________
Signature of researcher Date

This document has been printed on both sides of a single sheet of paper.
Please retain a copy of this document for your records.

Bibliography

J. Alderete. The syntax of Xiamen tone sandhi. *Phonology*, 18:201–253, 2001.

K. Ashley, L. Disch, D. Ford, E. MacSaveny, S. Parker, C. Unseth, A. Williams-Ngumbu, R. Wong, and B. Yoder. How many constraints are there? A preliminary inventory of OT phonological constraints. *Occasional Papers in Applied Linguistics*, 9:1–36, 2010.

M. Beckman and J. Edwards. Lengthenings and shortenings and the nature of prosodic constituency. In J. Kingston and M. Beckman, editors, *Papers in laboratory phonology 1: between the grammar and physics of speech*, pages 152–178. Cambridge University Press, Cambridge, 1990.

A. Belotel-Grenié and M. Grenié. The creaky voice phonation and the organization of Chinese discourse. In *Proceedings of the International Symposium on Tonal Aspects of Languages, With Emphasis on Tone Languages*, 2004.

L. Bloomfield. *Language*. Henry Holt and Company, New York, 1933.

Y.-C. Chang and Y.-C. Su. Tone modification of the third tone in Mandarin spoken in Taiwan. *Cahiers de Linguistique Asie Orientale*, 23:39–59, 1994.

Y. R. Chao. *Mandarin primer*. Harvard University Press, Cambridge, Mass., 1948.

J. Charles-Luce. The effects of semantic context on voicing neutralization. *Phonetica*, 50:140–149, 1993.

J. Charles-Luce. Cognitive factors involved in preserving a phonemic contrast. *Language and Speech*, 40:229–248, 1997.

J. Charles-Luce and D. A. Dinnsen. A reanalysis of Catalan devoicing. *Journal of Phonetics*, 15:187–190, 1987.

M. Chen. The syntax of Xiamen tone sandhi. *Phonology Yearbook*, 4:109–149, 1987.

M. Chen. *Tone sandhi: patterns across Chinese dialects*. Cambridge University Press, Cambridge, 2000.

T. Cho and P. Keating. Articulatory and acoustic studies of domain-initial strengthening in Korean. *Journal of Phonetics*, 29:155–190, 2001.

N. Chomsky and M. Halle. *Sound Patterns of English*. Harper Row, New York, 1968.

C.-T. Chuang, Y.-C. Chang, and F.-F. Hsieh. Productivity in Taiwanese tone sandhi redux. In W. Lee and E. Zee, editors, *Proceedings of the 17th International Congress of Phonetics Sciences*, Hong Kong, 2011. City University of Hong Kong.

A. Cohn. *Phonetic and phonological rules of nasalization*. PhD thesis, University of California, Los Angeles, 1990.

M. Di Paolo. Pronunciation and categorization in sound change. In K. Ferrara, B. Brown, K. Walters, and J. Baugh, editors, *Linguistic change and contact :Pro-*

ceedings of the 16th Annual Conference on New Ways of Analyzing Variation in Language, pages 84–92, Austin, 1988. University of Texas.

O. Dmitrieva, A. Jongman, and J. Sereno. Phonological neutralization by native and non-native speakers: The case of russian final devoicing. *Journal of Phonetics*, 38: 483–492, 2010.

A. Faber and M. Di Paolo. The discriminability of nearly merged sounds. *Language Variation and Change*, 7:35–78, 1995.

J. Fon. *What are tones really like? An acoustic-based study of Taiwan Mandarin tones.* PhD thesis, National Taiwan University, 1997.

M. Fourakis and G. K. Iverson. On the 'incomplete neutralization' of German final obstruents. *Phonetica*, 41:140–149, 1984.

G. R. Guy. Explanation in a variable phonology: An exponential model of morphological constraints. *Language Variation and Change*, 3:1–22, 1991.

S. Hargus and E. M. Kaisse. *Studies in Lexical Phonology.* Academic Press, San Diego, 1993.

J. Harris. *Phonological variation and change: studies in Hiberno-English.* Cambridge University Press, Cambridge, 1985.

L. Hwei-Bing. *Contextual Stability of Taiwanese Tones.* PhD thesis, University of Connecticut, Storrs, 1988.

S. Inkelas and C. Zoll. Is grammar dependence real? a comparison between cophonological and indexed constraint approaches to morphologically conditioned phonology. *Linguistics*, 45:133–171, 2007.

W. Jassem and L. Richter. Neutralization of voicing in Polish obstruents. *Journal of Phonetics*, 17:317–325, 1989.

P. Keating. Phonetic representation in a generative grammar. *Journal of Phonetics*, 18:321–334, 1990.

V. Kharlamov. Incomplete neutralization of the voicing contrast in word-final obstruents in Russian: Phonological, lexical, and methodological influences. *Journal of Phonetics*, 43:47–56, 2015.

P. Kiparsky. From cyclic phonology to lexical phonology. In H. v. d. Hulst and N. Smith, editors, *The Structure of Phonological Representations*, pages 131–175. Cambridge University Press, 1982.

P. Kiparsky. Some consequences of lexical phonology. In H. v. d. Hulst and N. Smith, editors, *The Structure of Phonological Representations*, pages 131–175. Cambridge University Press, 1985.

P. Kiparsky. Opacity and cyclicity. *The linguistic review*, 17:351–366, 2000.

R. Kirchener. Synchronic chain shifts in optimality theory. *Linguistic Inquiry*, 27: 341–350, 1996.

P. Kratochvil. *The Chinese Language Today: features of an emerging standard*. Hutchinson, London, 1968.

C.-H. Kuo. *Perception and acoustic correlates of the Taiwanese tone sandhi group*. PhD thesis, University of California, Los Angeles, 2013.

W. Labov. Methodology. In W. O. Dingwall, editor, *A survey of linguistic science*, pages 412–497, College Park, 1971. University of Maryland Linguistics Program.

W. Labov. On the use of the present to explain the past. In L. Heilmann, editor, *Proceedings of the 11th International Congress of Linguists*, pages 825–851, Bologna, 1975. Mulino.

W. Labov. *Principles of linguistic change*. Blackwell, Oxford and Cambridge, Mass., 1994.

W. Labov, M. Yaeger, and R. Steiner. *A quantitative study of sound change in progress*. U.S. Regional Survey, Philadelphia, 1972.

Y. Li. Tone sandhi and tonal coarticulation in Fuzhou Min. In *Proceedings of ICPhS 10*, Glasgow, 2015.

J. Lin. The syntax of Xiamen tone sandhi. *Phonology*, 11:237–275, 1994.

L. Lisker. "Voicing" in English: A catalogue of acoustic features signaling /b/ versus /p/ in trochees. *Language and Speech*, 29:3–11, 1985.

A. Manaster Ramer. Phonological and phonetic representation. *Journal of Phonetics*, 24:477–489, 1996.

J. Milroy and J. Harris. When is a merger not a merger? the meat/mate problem in a present-day English vernacular. *English World-Wide*, 1:199–210, 1980.

E. Moreton. Non-computable functions in Optimality Theory. In J. J. McCarthy, editor, *Optimality Theory in Phonology*, pages 141–164. 2004.

E. Moreton and P. Smolensky. Typological consequences of local constraint conjunction. In L. Mikkelsen and C. Potts, editors, *WCCFL 21 Proceedings*, Cambridge, MA, 2002. Cascadilla Press.

D. Mortensen. Semper Infidelis: Theoretical dimensions of tone sandhi chains in Jingpho and AHmao. Unpublished, 2002.

J. Myers and J. Tsay. Investigating the phonetics of Mandarin tone sandhi. *Taiwan Journal of Linguistics*, 1:29–68, 2003.

J. Myers and J. Tsay. Neutralization in Taiwan Southern Min tone sandhi. In Y. E. Hsiao, H.-C. Hsu, L.-H. Wee, and D.-A. Ho, editors, *Interfaces in Chinese phonology: Festschrift in honor of Matthew Y. Chen on his 70th birthday*, pages 47–78. Academia Sinica, 2008.

G. Nunberg. A falsely reported merger in eighteenth-century English: a study in diachronic variation. In W. Labov, editor, *Locating language in time and space*, pages 221–250, New York, 1980. Academic Press.

S.-H. Peng. Production and perception of Taiwanese tones in different tonal and prosodic contexts. *Journal of Phonetics*, 25:371–400, 1997.

S.-H. Peng. Lexical versus 'phonological' representations of Mandarin sandhi tones. In B. M. and P. J, editors, *Papers in laboratory phonology 5: acquisition and the lexicon*, pages 152–167. Cambridge University Press, 2000.

J. Pierrehumbert. Phonological and phonetic representation. *Journal of Phonetics*, 18:375–394, 1990.

R. F. Port and P. Crawford. Incomplete neutralization and pragmatics in German. *Journal of Phonetics*, 17:257–282, 1989.

R. F. Port and M. L. O'Dell. Neutralization of syllable-final voicing in German. *Journal of Phonetics*, 13:455–471, 1985.

P. Price, M. Ostendorf, S. Shattuck Hufnagel, and C. Fong. The use of prosody in syntactic disambiguation. *Journal of the Acoustical Society of America*, 90:2956–2970, 1991.

A. Prince and P. Smolensky. *Optimality Theory: constraint interaction in generative grammar*. Blackwell, Malden, 2004. Available on ROA (http://roa.rutgers.edu).

J. Rubach. Final devoicing and cyclic syllabification in German. *Linguistic Inquiry*, 21:79–93, 1990.

X. N. Shen. A pilot study on the relation between the temporal and syntactic structures in Mandarin. *Journal of the International Phonetic Association*, 22:35–43, 1992.

L. Slowiaczek and D. Dinnsen. On the neutralizing status of Polish word-final devoicing. *Journal of Phonetics*, 14(3):325–341, 1985.

L. Slowiaczek and H. Szymanska. Perception of word-final devoicing in Polish. *Journal of Phonetics*, 17(3):205–212, 1989.

D. Steriade. Paradigm uniformity and the phonetics-phonology boundary. In M. Broe and J. Pierrehumbert, editors, *Papers in Laboratory Phonology*, Cambridge, 2000. Cambridge University Press.

P. Trudgill. *The social differentiation of English in Norwich*. Cambridge University Press, Cambridge, 1974.

A. Turk and S. Shattuck-Hufnagel. Segmental durations in the vicinity of prosodic phrase boundaries. *Phonology*, 35:445–472, 2007.

B. van Rooy, D. Wissing, and D. D. Paschall. Demystifying incomplete neutralization during final devoicing. *Southern African Linguistics and Applied Language Studies*, 21:49–66, 2003.

W. Wang. Phonological features of tone. *International Journal of American Linguistics*, 33(2):93–105, 1967.

N. Warner, A. Jongman, J. Sereno, and R. Kemps. Incomplete neutralization and other subphonemic durational differences in production and perception: Evidence from Dutch. *Journal of Phonetics*, 32:251–276, 2004.

C. Wightman, P. Price, M. Ostendorf, and S. Shattuck-Hufnagel. Segmental durations in the vicinity of prosodic phrase boundaries. *Journal of the Acoustical Society of America*, 91:1707–1717, 1992.

Y. Xu. Contextual tonal variations in Mandarin. *Journal of Phonetics*, 25:61–83, 1997.

M. Yip. *Tone*. Cambridge University Press, Cambridge, 2002.

A. C. Yu. Understanding near mergers: The case of morphological tone in Cantonese. *Phonology*, 24:187–214, 2007.

E. Zee. A spectrographic investigation of Mandarin tone sandhi. *UCLA Working Papers*, 49:98–116, 1980.

J. Zhang, Y. Lai, and C. Turnbull-Sailor. Opacity, phonetics, and frequency in Taiwanese tone sandhi. In *Current Issues in Unity and Diversity of Languages: Collection of Papers Selected from the 18th International Congress of Linguists*, pages 3019–3038. Linguistic Society of Korea, 2009.

B. Zhiming. Chinese tone sandhi. In M. van Oostendorp, C. Ewen, E. Hume, and K. Rice, editors, *The Blackwell companion to phonology*, volume 5, pages 2561–2585. Wiley-Blackwell, 2011.

E. C. Zsiga. *Features, gestures, and the temporal aspects of phonological organization.* PhD thesis, Yale University, 1993.

E. C. Zsiga. An acoustic and electropalatographic study of lexical and post-lexical palatalization in American English. In B. Connell and A. Arvaniti, editors, *Phonology and phonetic evidence, papers in laboratory phonology IV*, pages 282–302, Cambridge, 1995. Cambridge University Press.

E. C. Zsiga. Features, gestures, and Igbo vowels: an approach to the phonology/phonetics interface. *Language*, 73:227–274, 1997.