COORDINATION AND INTERACTION

IN MARKEDNESS SUPPRESSION

by

George-Michael Pescaru

A thesis submitted to the faculty of
The University of Utah
in partial fulfillment of the requirements for the degree of

Master of Arts

Department of Linguistics

The University of Utah

May 2015

The University of Utah Graduate School

STATEMENT OF THESIS APPROVAL

The thesis of **George-Michael Pescaru**

has been approved by the following supervisory committee members:

Aaron Kaplan	, Chair	**9/4/2014** Date Approved
Abby Kaplan	, Member	**9/4/2014** Date Approved
Edward J. Rubin	, Member	**9/4/2014** Date Approved
Shannon Barrios	, Member	**9/4/2014** Date Approved

and by **Edward J. Rubin** , Chair/Dean of

the Department/College/School of **Linguistics**

and by David B. Kieda, Dean of The Graduate School.

ABSTRACT

Markedness Suppression (MS) is a formal treatment of gradient variation in Optimality Theory that permits certain markedness constraints to have any number of violations "suppressed," or ignored, by EVAL. This theory was designed for systems where the "locus of variation" is a single segment/feature whose variability has little impact on the well-formedness of the rest of the output ("local" variation), as is the case for French schwa deletion. However, it is not immediately clear how MS can account for some other kinds of variation: in particular, "coordinated" variation, variation where the choice to suppress at one locus is conditioned by other suppression decisions, and variation at a single locus that is differently conditioned by "interaction" with neighboring phonological environments. These types of variation conflict, respectively, with the assumption of independence between loci and with the inability of suppressible constraints to reliably dominate the constraints ranked below them.

This thesis analyzes two coordinated patterns – free variation between all-[p] and all-[b] in Warao and Shimakonde midvowel reduction, which varies in the extent of its application – and two interactive patterns – vowel backness harmony in Hungarian and English t/d-deletion. I will show that although suppressible constraints overgenerate on their own in the case of coordination, we can account for these patterns by introducing additional constraints that handle the additional structural requirements of variation. For the case of interaction, although MS does not permit the full range of possible analyses

because of its weaker constraint dominance, I will show that it is possible to produce analyses of these patterns so long as all suppressible constraints refer to particular environments internally and a stringency relationship exists between them.

In showing that analyses of these systems are possible, this thesis demonstrates that MS can account for these classes of variation and, thus, is not limited as a theory of solely local variation.

TABLE OF CONTENTS

LIST OF TABLES

ACKNOWLEDGEMENTS

I owe an enormous amount of thanks and credit to my advisory committee – Aaron Kaplan, Abby Kaplan, Ed Rubin, and Shannon Barrios – for their help making this thesis a reality, as well as for bending over backward to help me meet various deadlines. I especially owe thanks to Aaron Kaplan, who spent numerous hours pointing out where, exactly, I had gone horribly wrong and whose regular commentary and critiques were indispensible in the shaping of this thesis over the course of seven drafts.

I also owe special thanks to my fellow students and the department of linguistics for being an awesome bunch and to my family for being my family. I really appreciate that.

CHAPTER 1

INTRODUCTION

1.1 Variation in phonology

All speech is, to some degree, variable on a PHONETIC level. Such variation is gradient, occurring within a single phonological category. For example, the precise formant frequency associated with a given vowel (e.g. [i]) differs both between speakers and between utterances for a single speaker. Many languages also have variable PHONOLOGICAL patterns in which mutually exclusive output patterns are synchronically attested, often in the same speaker. For example, in French, [ə] can be optionally deleted in most circumstances, leading to multiple possible realizations of a single word or phrase, each with different degrees of [ə] deletion (Dell 1973, Côté 2000, and numerous others).

(1) *envie de te le demander* 'feel like asking you for it'

a. ãvidətələdəmãde
b. ãvid_tələdəmãde
c. ãvidət_lədəmãde
d. ãvidətəl_dəmãde
e. ãvidətələd_mãde
f. ãvid_təl_dəmãde
g. ãvid_tələd_mãde

(Côté 2000)

In (1), deleted [ə] are represented with underscores (_). As shown in (1a), the phrase *envie de te le demander* has at most four [ə] – which I will refer to as LOCI OF VARIATION, meaning positions where some feature, segment, or morpheme may be optionally realized – but it may also be realized with as few as two (in (1f,g)), and there are seven possible combinations of deletion and/or preservation possible in all. Though some of these possible outputs may be more frequently attested than others, all of them are grammatical realizations of (1).

Variation in (1) is LOCAL (Kimper 2011a) in the sense that the choice of output pattern (i.e. delete or not) is made independently at each locus of variation, but this is not the only possibility. Warao, for example, has a pattern of GLOBAL variation (Kimper 2011a) where an individual word can be realized with either [p] or [b], but cannot have a mixture of the two.

(2) /paro + parera/ ‘weak’

a. **p**aro**p**arera
b. **b**aro**b**arera
c. ***p**aro**b**arera
d. ***b**aro**p**arera

(Osborn 1966)

In (2), the choice between output patterns at one locus essentially "locks in" the choice for all other loci. A similar pattern of interdependence among loci is found in Shimakonde. Vowel reduction of pretonic (= prepenultimate) midvowels to [a] in Shimakonde is subject to a set of requirements Liphola (2001) terms CONTINUITY OF REDUCTION (CORE), which requires reduction to begin at the left edge and not skip any midvowels along the way (in the data in (3), tone is marked with diacritics, while stress is marked by its IPA symbol (ˈ)).

(3) kú-pélévéléˈlééla ‘to not reach a full size for’

a. kú-pélévéléˈlééla
b. kú-p**á**lévéléˈlééla
c. kú-p**á**l**á**véléˈlééla
d. kú-p**á**l**á**v**á**léˈlééla
e. kú-p**á**l**á**v**á**l**á**ˈlééla
f. *kú-pél**á**v**á**l**á**ˈlééla
g. *kú-pélév**á**l**á**ˈlééla
h. *kú-pélévél**á**ˈlééla
i. *kú-p**á**lév**á**léˈlééla

(Liphola 2001)

In contrast to French and Warao, where the loci of variation were unified only by their shared features, the variation in (3) is in the extent to which the reduction process occurs, if it occurs at all. However, due to CORE, the variation at each individual locus is similar to the global variation in Warao: the available options at each locus, except for the leftmost, depend on the choices made at the loci to the left. Additionally, because variation is possible whenever certain features are present in an individual segment, each of these three patterns can have arbitrarily many loci of variation.

Another class of variable patterns has its loci CREATED, rather than nullified, by particular structural configurations. Consequently, these patterns have, at most, one locus of variation.[1] An example of this is vowel backness harmony in Hungarian, where the pairing of harmonically-neutral front vowels, like [e:], with harmonic back vowels, like [ɔ**]**, are variably transparent with respect to suffix vowel harmony (relevant vowels in bold).

In (4), the combination of vowels [ɔ ... e:] in the stem includes a harmonic back vowel, [ɔ], as well as a harmonically-neutral front vowel, [e:]. Stems with this

[1] This is the case if variation is interpreted as occurring at the level of a word, following Shimakonde and Warao, and not the phrase level, following French. Given multiple words, there could be multiple loci of variation.

configuration of vowels can be either transparent with respect to suffix vowel harmony, selecting [+back] suffix vowels as in (4a), or they can be opaque, selecting [-back] suffix vowels, as in (4b). Just as in the preceding examples of harmony, both of these realizations are grammatical.

(4) Variable harmony with harmonic + neutral stems

a. [**ɔ**rze:n-n**ɔ**k] ***a****rzénn****a****k* 'arsenic-DAT'
b. [**ɔ**rze:n-n**ɛ**k] ***a****rzénn****e****k* 'arsenic-DAT'

(Hayes & Londe 2006)

In (4), the combination of vowels [ɔ ... e:] in the stem includes a harmonic back vowel, [ɔ], as well as a harmonically-neutral front vowel, [e:]. Stems with this configuration of vowels can be either transparent with respect to suffix vowel harmony, selecting [+back] suffix vowels as in (4a), or they can be opaque, selecting [-back] suffix vowels, as in (4b). Just as in the preceding examples of harmony, both of these realizations are grammatical.

In all of the examples of variable patterns given so far, certain variants are more common than others. For the first class of patterns (i.e. French, Warao, and Shimakonde), however, this can be attributed both to the structural configuration of each locus as well as the number of loci, which largely determines how many variants are possible (and, consequently, has an effect on the frequencies for each variant). For variation in the style of Hungarian, differences in frequency are conditioned almost exclusively by the composition of each of the possible structural configurations, as there is only one word-internal locus of variation. In Hungarian, for instance, some harmonically neutral front vowels select [-back] suffix vowel variants more often than others, while stems with MULTIPLE harmonically neutral front vowels are more likely to choose [-back] suffix vowels no matter what vowels are involved.

Similarly, in most (if not all) English dialects, [t d] in a word-final consonant cluster (e.g. *wes**t***) is likely to be deleted, but the frequency of deletion depends on the following segment, as indicated in the Table 1.1.

Additionally, deletion is sensitive to the morphological content of the [t] or [d]: for example, the [t] of a past tense morpheme is less likely to be deleted than a [t] that is merely at the end of a stem morpheme.

A number of theoretical approaches within and without OPTIMALITY THEORY (Prince & Smolensky 1993) have been devised to handle gradient variation, beginning with the PARTIAL ORDERS theory (Antilla 1997, 2007, Kiparsky 1993), where variation was achieved by reranking constraints to achieve different outcomes.

Subsequently, new theories have been proposed to address the weaknesses of Partial Orders. MARKEDNESS SUPPRESSION (Kaplan 2008, 2011) is one of these, having been developed to handle "local" variation of the kind of French, where variation at each locus is essentially independent of all others. The purpose of this thesis is to show that Markedness Suppression can also handle the other kinds of variation listed above, contrary to some objections to the theory (see especially Kimper 2011a).

The following section will present the theory of Markedness Suppression.[2] Section 1.3 will then give a brief summary of the thesis, breaking down the problems presented by Warao, Shimakonde, Hungarian, and English into two broader categories of variable patterns – those that show coordination between loci and those that show interaction with environments – and explaining why they might pose problems for

[2] In Chapter 2, I will provide a brief overview of three other theories of variation in order to give some background to Markedness Suppression: Partial Orders (Antilla 1997, 2007, Kiparsky 1993), Rank-Ordered EVAL (Coetzee 2004, 2006), and Stochastic Optimality Theory (Boersma & Hayes 2001).

TABLE 1.1. t/d Deletion in African American English (Washington, D.C.) (Labov, Cohen, Robbins, & Lewis 1968)

Context	Frequency of [t d] deletion
wɛs[t]# C ~ wɛs# C	76%
wɛs[t]# V ~ wɛs# V	29%
wɛs[t]# # ~ wɛs# #	73%

Markedness Suppression as well as how I will approach analyses to these patterns to solve those problems. These analyses will provide a basis for generalizing to other systems whose patterns run into one or both of these problems.

1.2 Markedness Suppression

As noted above, Markedness Suppression is one of a number of formal approaches to variation in Optimality Theory. Under Markedness Suppression, the variability of output is expressed as the variable parsing of violations by EVAL: essentially, certain markedness constraints[3] – those motivating a particular variable alternation, denoted formally with an optionality operator (Θ) – permit any number of violations on a per-candidate basis to be SUPPRESSED, or thrown out, permitting candidates that would otherwise have been eliminated by that constraint to be selected as optimal.

Each decision to suppress a violation at a particular locus is random, with probability p – the implication being that the choice at one locus is independent of all other choices to suppress (such that the probability of n violations being suppressed for a given candidate is p^n). The number p is set on a per-language basis – which means that a

[3] Faithfulness constraints are excluded, as suppressing their violations would predict potentially pathological behavior. See Kaplan (2011: 338) for discussion.

language with multiple variable constraints would be predicted to suppress violations from each of those constraints at the same rate – and *p* is meant to "summarize" all of the various influences on the rate of variation (e.g. speech rate, formality), and in practice its value is determined in such a way as to approximate the frequencies of each output pattern in some corpus. Manipulation of this parameter allows analyses under Markedness Suppression to predict how certain output patterns will be more frequent than others, reflecting persistent empirical differences in the proportional frequency of variable patterns.

A typical example of how Markedness Suppression functions is variable [ə] deletion in French. As described above, deletion of [ə] is permitted, variably, in most circumstances, but is not possible when it would create certain illicit consonant clusters (Côté 2000). In Table 1.2, I reproduce the data in (1) along with several ungrammatical examples (cf. Côté 2000: example 44).

In Table 1.2, deleted are represented with underscores (_).Examples (b-e) show that deletion of a single schwa, creating a cluster of two consonants, is always acceptable.

TABLE 1.2. Realizations of *envie de te le demander* ‘feel like asking you for it’

Attested	Unattested
a. ãvidətələdəmãde	k. *ãvidət_l_dəmãde
b. ãvid_tələdəmãde	l. *ãvid_t_l_dəmãde
c. ãvidət_lədəmãde	m. *ãvid_t_l_d_mãde
d. ãvidətəl_dəmãde	n. *ãvid_t_l_d_mãde
e. ãvidətələd_mãde	o. ?ãvidətəl_d_mãde
f. ãvid_təl_dəmãde	p. ?ãvid_təl_d_mãde
g. ãvid_tələd_mãde	
h. ãvidət_ləd_mãde	
i. ãvid_t_lədəmãde	
j. ãvid_t_ləd_mãde	

Deletion of multiple schwas is acceptable in some cases (f-j), categorically unacceptable in others (k-n), and marginal in two cases, (o) and (p). Leaving the last two cases aside here (see Côté 2000 and Kaplan 2011 for a full discussion), the primary difference between the deletion patterns in (f-j) and those in (k-n) is that deletion in the latter cases creates a three-consonant cluster where the central consonant is the most sonorous (namely [tld]), but no such cluster is created in (f-j). The tableau in (5) (adapted from Kaplan 2011) illustrates how Markedness Suppression can account for this, using the constraint *CNC to ban clusters like [tld].

(5) Tableau for *envie de te le demander* 'feeling like asking you for it'

/ãvi də tə lə dəmãde/	*CNC	ʘ *ə	MAX
(→) a. ãvi də tə lə dəmãde		****	
(→) b. ãvi d_ tə lə dəmãde		***	*
(→) c. ãvi də t_ lə dəmãde		***	*
(→) d. ãvi də tə l_ dəmãde		***	*
(→) e. ãvi də tə lə d_mãde		***	*
→ f. ãvi d_ tə l_ dəmãde		**	**
→ g. ãvi də t_ lə d_mãde		**	**
h. ãvi də t_ l_ dəmãde	*!	**	**

Here, variation is motivated by ʘ *ə. The only candidate eliminated by *CNC, the constraint used to eliminate the illicit consonant clusters mentioned above, is (h). All of the other candidates are possible outputs, which I mark throughout this thesis with a parenthetical arrow, (→). Which candidate is determined to be optimal depends on how violations of ʘ *ə are suppressed. For example, (a) wins if at least two of its violations of ʘ *ə are suppressed while no other candidates have violations suppressed. We can represent violations that have been suppressed by using *o* instead of * for those violations, as in (6).

(6) Tableau for *envie de te le demander* 'feeling like asking you for it' ((a) is the winner)

/ãvi də tə lə dəmãde/	*CNC	Θ *ə	MAX
→ a. ãvi də tə lə dəmãde		oo**	
b. ãvi d_ tə lə dəmãde		***!	*
c. ãvi də t_ lə dəmãde		***!	*
d. ãvi də tə l_ dəmãde		***!	*
e. ãvi də tə lə d_mãde		***!	*
f. ãvi d_ tə l_ dəmãde		**	*!*
g. ãvi də t_ lə d_mãde		**	*!*
h. ãvi də t_ l_ dəmãde	*!	**	**

With two violations of Θ *ə suppressed for (a) alone, and with no other violations suppressed, candidates (b-e) are eliminated by that constraint for having one extra violation beyond the minimum among remaining candidates. (a) prevails over the remaining candidates, (f) and (g), because it has no violations of MAX, where the other two have two each. Note that no candidates in (6) are marked with (→), as this tableau ignores the possibility of other violations being suppressed in order to show the results of a particular pattern of suppression.

Each of the variants in (5) occurs at a different frequency (Kaplan 2011). As discussed above, Markedness Suppression assigns each language a probability of suppression, p, which is the rate at which any given violation is suppressed. Since the rate of suppression is statistically independent for each violation, the rate at which n violations are suppressed is p^n. So, for example, if $p = 0.5$ for French, then the probability of both violations of Θ *ə being suppressed for (5g) is $p^2 = 0.25$. (This is NOT, however, the probability of the output pattern in (5g) because some of the violations of other candidates must be RETAINED in order for (5g) to win with two of its violations suppressed.)

Similarly, the probability that only a subset of violations are suppressed is equal to the PRODUCT of the probability of all of the relevant violations being suppressed AND the probability of all of the relevant violations NOT being suppressed – which, for *m* violations RETAINED, is $(1 - p)^m$. For example, the probability that EXACTLY two of (5a)'s four violations of ʘ *ə will be suppressed is $p^2(1 - p)^2$ – which, if $p = 0.5$, is $(0.5)^2(1 - 0.5)^2 = 0.5^4 = 0.0625$.

If a given candidate requires *n* of its own violations to be suppressed and *m* other violations to be retained in order to win, then, following the above, its probability of winning is $p^n(1 - p^m)$. If a candidate wins under multiple different scenarios of suppression and retention of violations, then its probability of winning is the sum of the probabilities of each of those scenarios. For French, these calculations become rather involved; again, see Kaplan (2011) for a full discussion.

Since Markedness Suppression is, as discussed here, capable of modeling the frequency with which a given output pattern occurs, where data are available, I will compare the absolute frequency predictions made by the analysis to the attested frequencies of a corpus. However, it is arguable that the ability to model attested frequencies in absolute terms is not as important as the ability to accurately model the RELATIVE frequencies of each output form (as well as whether a phenomenon is categorical, e.g. categorical deletion in some environment)(see Coetzee 2004), so I will focus primarily on the relative frequencies predicted by each analysis.

So far, I have shown how Markedness Suppression works. The next section, I will take another look at the data given in Section 1.1 and show how those data highlight some limitations to Markedness Suppression, as well as how I intend to overcome those limitations.

1.3 The problems of coordination and interaction

Markedness Suppression was intended primarily to account for patterns of variation like that of French [ə] deletion, where variation is local, or decided independently at each of potentially many loci of variation. The other patterns of variation discussed above all deviate from this description of variation in one of two ways.

Warao and Shimakonde demonstrate coordination between loci: in particular, the variability of an output form can "turn off" at some loci given choices made at other loci. In formal terms, this seems to violate the assumption by Markedness Suppression that the decision to suppress or retain a violation is independent of the same decisions made for other violations. I will refer to this as the COORDINATION PROBLEM.

Hungarian and English, on the other hand, show variation in particular structural configurations, or environments, where the frequency of certain variants differs on the environment. Formally, this requires the use of different constraints to account for the influences of different environments – but since suppressible constraints cannot reliably dominate other constraints (as all their violations can be suppressed), analyses of this phenomenon cannot rely on constraint domination to achieve the correct output patterns and frequency relations between them. I will refer to this as the INTERACTION PROBLEM.

The following two sections will give a more detailed review of the formal problems posed by these patterns as well as how I intend to avoid them. Additionally, the formal strategies by which these problems are solved make predictions as to the kinds of "coordinated" and "interactive" variations that are possible, due to the assumptions underlying these solutions. In particular, both patterns of coordination as predicted to pattern according to some phonetic or phonological relationship between each variant or context. These predictions are given in more detail below.

1.3.1 Coordination

As noted above, both Warao and Shimakonde have variable patterns that could be construed as coordinated variation. The Warao pattern of free variation between [p] and [b], described in Section 1.1, demonstrates a subset of the coordination problem that Kimper (2011a) describes as "global" variation. As opposed to local variation as in the French schwa pattern, where variation occurs freely at each locus, here variation appears to be the result of a setting that applies throughout the word (i.e. choose either [p] or [b]), as shown in the data reproduced below from (2).

(7) /paro + parera/ ‘weak’

a. **p**aro**p**arera
b. **b**aro**b**arera
c. ***p**aro**b**arera
d. ***b**aro**p**arera

(Osborn 1966)

Shimakonde vowel reduction shows another kind of coordinated variation. As described in Section 1.1, in Shimakonde, prepenultimate (= pretonic) midvowels optionally reduce, but reduction must proceed from the left edge without skipping any midvowels along the way, as shown in (8) (reproduced from (3)).

(8) kú-pélévéléˈlééla ‘to not reach a full size for’

a. kú-pélévéléˈlééla
b. kú-p**á**lévéléˈlééla
c. kú-p**á**l**á**véléˈlééla
d. kú-p**á**l**á**v**á**léˈlééla
e. kú-p**á**l**á**v**á**l**á**ˈlééla
f. *kú-pél**á**v**á**l**á**ˈlééla

g. *kú-pélév**á**l**á**ˈlééla
h. *kú-pélévél**á**ˈlééla
i. *kú-p**á**lév**á**léˈlééla

(Liphola 2001)

This is a coordinated pattern insofar as reduction is blocked for the remainder of the word after any midvowel fails to reduce; a single locus has influence on the possibilities at other loci.

Markedness Suppression is expected to have some difficulty with these patterns because suppression at one locus cannot be conditioned on suppression elsewhere. Consequently, this is an area in which Markedness Suppression may overgenerate, predicting all of the logically possible combinations of variants. For example, for Warao, a suppressible constraint meant to account for the variation between [p] and [b], Θ *P – which favors [b] when violations are retained and [p] when violations are suppressed – would permit not only words with all [p] or all [b], but all combinations of the two, just as the analysis of French in Section 1.2 permitted multiple combinations of deletion and retention of [ə]. This is shown in (9).

(9) Overgeneration with a suppressible constraint in Warao

/paro + parera/	Θ *P	*OBSVOICE
(→) a. paroparera	**	
(→) b. barobarera		**
💣 c. parobarera	o	*
💣 d. baroparera	o	*

The key fact is that not even the French data permitted ALL combinations of deletion and retention in every position, and the method of addressing that problem used there – adding a constraint that penalizes certain illicit structural configurations that could arise from [ə] deletion – will also be suitable for analyzing the more obvious coordination of Warao and Shimakonde. In both languages, the suppressible constraint in each case can be chosen to decide between only the features that are variable (e.g. ⊙ *P for variation between [p] and [b]), while other constraints can enforce the remaining structural requirements (e.g. voicing harmony to prevent combinations of [p] and [b]). This predicts that the coordination in a variable process will not be random, but instead will result from the structural requirements imposed by some other attested process.

The coordination problem will be addressed in full detail in Chapter 3.

1.3.2 Interaction

As noted above, Hungarian and English each have variable patterns where the variable alternation between features (or deletion) is the same in each environment, but where each environment has a different influence on the frequency of a particular variant. In Hungarian, harmonically neutral front vowels engender variable transparency in harmony when paired with harmonic back vowels, as shown in (10) (reproduced from (4)).

(10) Variable harmony with harmonic + neutral stems

a. [**ɔ**rze:n-n**ɔ**k]	***a****rzénn****a****k*	'arsenic-DAT'
b. [**ɔ**rze:n-n**ɛ**k]	***a****rzénn****e****k*	'arsenic-DAT'

(Hayes & Londe 2006)

In many dialects of English, [t] and [d] optionally delete word-finally when following another consonant, but the frequency of deletion depends on the following segment (and the morphological character of the [t] or [d]). An example of this from one dialect of English is given in Table 1.1.

These patterns provide potential examples of interaction due to the multitude of elements conditioning variation. For Hungarian, as discussed in Section 1.1, each of the harmonically neutral vowels trigger harmony at different rates, and the number of neutral vowels is also a factor. For English, as just stated, both the following segment – a consonant, vowel, or pause – and the morphological environment condition deletion.

Whether Markedness Suppression has difficulty with a pattern of this kind depends almost entirely on the constraints used to describe it. To begin with, any analysis of these patterns will require multiple constraints, each referring to different environments (or classes of environments). From there, there are two general obstacles to overcome.

First, because the probability of suppression p is assumed to be determined on a per-language basis, the only way to account for systemic differences in the frequency of each pattern is to use constraints that in some way assign more violations to the less-frequent variants. Due to the way Markedness Suppression accounts for suppression of multiple violations (Section 1.2), this makes the not necessarily valid prediction that differences in frequencies between variants will follow an exponential pattern; that is, the most common variant has probability p, the next p^2, and so on. More problematically, multiple aspects of a given structural configuration can impact frequencies simultaneously. In Hungarian, for instance, stems with two neutral vowels are uniformly more likely to select [-back] suffix vowels than those stems with one neutral vowel, regardless of which vowels are involved, but the kind of neutral vowels still influences probabilities within the class of stems with two. Whether the analysis can account for this nicely depends to a large extent on what constraints are used.

The second difficulty arises in selecting constraints capable of generating the correct set of optimal outputs. Suppressible constraints can never be assured to actively dominate the constraints they outrank: all of their violations can be suppressed, effecting a state in which the suppressible constraint has no impact on EVAL at all. As such, Markedness Suppression cannot rely on the consequences of constraint interaction to do any work toward a particular outcome when it comes to the relationship between a suppressible constraint and the constraints it dominates.

Consequently, systems that REQUIRE precise hierarchical relationships between constraints may not be amenable to a Markedness Suppression analysis. For example, consider a language where word-final [t] is always preserved, but word-final [d] is sometimes deleted. If the constraints used to describe this are FINAL-C (McCarthy 1993), which requires word-final consonants, plus *T, and *D, then it must be the case that **Θ** FINAL-C >> *D, with FINAL-C suppressible, in order to permit [d] to delete some of the time, but it must ALSO be the case that FINAL-C >> *T, with FINAL-C *not* suppressible, in order to categorically prevent [t] from deleting. The ranking Θ FINAL-C >> *D, *T permits BOTH [t] and [d] to be deleted, contrary to the specifics of the pattern.

Formally stated, suppose there is a system with constraints Θ X, Y, and Z, where Θ X must outrank Y and Z, as in (11).

(11) Θ X >> Z, Y

Now suppose that, in order to account for the data, violations of Θ X must be suppressible such that Z can decide between certain candidates, but never Y. That is, Θ X is variable with respect to Z, but not Y. Under Markedness Suppression, there is no way to achieve the desired result from the configuration in (11). Θ X is variable with respect to Z, as desired, but it must also be variable with respect to Y, which we do not want; there is no way to distinguish its relationship to Z from its relationship to Y so long as it outranks both. Variation thus cannot be limited to the ranking between a subset of constraints.

Patterns born out of interactions between environments, like Hungarian vowel harmony, have been likely to suggest analyses of this kind. For example, one generalization of Hungarian vowel harmony suggests that back harmony occurs when the last neutral vowel is [i] but not when it is [e:]. This yields the structure in (11) if we undertake the analysis that the constraint motivating harmony, Θ X, dominates two constraints penalizing it in each environment: Z penalizing harmony through [e:] and Y through [i], for instance.[4] To a large extent, the association between the structure in (11) and this class of variable patterns is due to the theoretical basis of the analyses previously undertaken for these patterns, which relied on constraint domination. As I will show with Hungarian and English, the structure in (11) is not a necessary result of an analysis of this class of pattern.

In order to work around these two obstacles – getting the frequencies right and getting the right output patterns – there are two requirements placed on the analysis. First, those candidates that are systematically less frequent (due to the environments involved) must always receive more violations total than those that are systematically more frequent. This will be achieved with STRINGENCY (de Lacy 2003), which requires there to be a natural hierarchy relating the relevant environments (e.g. vowel height in Hungarian). Consequently, since this a Markedness Suppression analysis of these phenomenon requires stringency, this predicts that ALL variable patterns that show differentiation in frequencies depending on the environment are the result of an

[4] This example was provided by W. Kimper (personal communication, Nov. 19, 2013). A variant of this, based on Kimper (2011b), will be discussed in more detail in Chapter 4.

interaction between a variable pattern and a natural hierarchy.

Second, relations of domination between multiple constraints must be consistent with the reality that high-ranking suppressible constraints cannot be relied on to enforce their violations over the constraints they dominate (following the discussion around (4) above). When environments interact, as they do in these languages, this means that each constraint must be specific enough to be able to handle everything pertaining to a given environment; it cannot dump off anything onto the ranking system. It should not be necessary for suppressible constraints to be ranked crucially with respect to one another. As I will show in analyzing these languages, these are not unreasonable requirements to meet.

The interaction problem will be addressed in full detail in Chapter 4.

CHAPTER 2

OTHER THEORIES OF VARIATION

2.1 Introduction

Markedness Suppression is just one of many theories of variation. Among the earliest and most influential theories is the Partial Orders theory (Antilla 1997, 2007, Kiparsky 1993), which treats variation as the result of the absence of a total order on CON so that the ranking between two unordered constraints is adopted on the fly and may differ from evaluation to evaluation. STOCHASTIC OPTIMALITY THEORY (Boersma & Hayes 2001) is another theory that achieves variation through the reranking of constraints, except unlike the Partial Orders theory, the rankings here are stochastic, such that the relative ranking of each individual pair of constraints is associated to some probability (not necessarily 1). Some theories (including Markedness Suppression) make more significant departures from Partial Orders. Among the more prominent of these are RANK-ORDERED EVAL (Coetzee 2004, 2006), in which only a subset of constraints can actually eliminate candidates, while the other constraints simply make the candidates they penalize less frequent winners, as well as HARMONIC SERIAL VARIATION (Kimper 2011a,b), based on the nonparallel HARMONIC SERIALISM (McCarthy 2000), an alternative to OT, in which candidates make multiple passes through EVAL with the winner at each stage becoming the input for the next stage and candidates differing by at most one change per stage, up until the input for some stage is also the optimal candidate at that stage, at which point the derivation is said to converge. Variation in Harmonic Serial Variation is achieved through the Partial Orders theory, permitting different rankings at different stages of the cycle.

The goal of this chapter is not to provide an assessment of these theories, but rather, to provide some background for the theory of Markedness Suppression and to situate it in the context of competing theories of variation. Since Chapter 4 will either use constraints taken from Partial Orders, Stochasic OT, or Rank-Ordered EVAL analyses of Hungarian and English or compare Markedness Suppression to one of those theories, I will focus on those here, providing illustrations of how each of them work through an analysis of French schwa deletion. Section 2.2 will discuss both Partial Orders and Stochastic OT, which are very closely related. Section 2.3 will discuss Rank-Ordered EVAL.

2.2 Variation through variable rankings

The theory that variation can be accounted for through the use of multiple constraint rankings begins with Kiparsky (1993), who uses this approach to analyze English t/d deletion.[5] Antilla (1997, 2007) provides most of the theoretical foundation for the theory, including a method for modeling the frequencies of particular variants on the basis of how many rankings cause those variants to win.

The basic mechanism for this theory is simple. Prince and Smolensky (1993) originally conceived of CON as a total order on a set of constraints. Under the Partial Orders theory (Antilla 1997, 2007) and very similarly for its derivatives, this is not

[5] The details of this analysis are given in Chapter 4.

necessarily the case: a subset of constraints might have no fixed ordering with respect to one another. Faced with two constraints with no ranking between them, EVAL will randomly impose a ranking between them. Thus, constraints X and Y, which have no ranking in CON, will be ranked X >> Y in some evaluations and Y >> X in others. This results in variation if at least two of the possible rankings differ in the output patterns they favor.

For the Partial Orders theory, there is a 50-50 chance of either ranking. Consequently, the frequency of a particular variant is determined by the proportion of possible rankings that favor it. For example, there are six possible rankings between the constraints X, Y, and Z. If three of those rankings favor candidate A, that candidate is predicted to win 50% of the time; if two of them favor candidate B, it is expected to win approximately 33.3% of the time; and if only one of them favors candidate C, it will win approximately 16.5% of the time.

The most significant change Stochastic OT makes to this paradigm is to permit one ranking to be favored over another: for X >> Y to happen 70% of the time, say, and Y >> X only 30% of the time. This means that in order to determine model frequencies for variants under Stochastic OT, it is not sufficient to count up the number of rankings that favors those variants; it is also necessary to take into account how often those rankings occur, given the possibility of favoritism. In order to achieve this random variation, Stochastic OT assumes that the ranking of each constraint in CON is associated with some amount of noise, which may cause constraints to randomly assume other rankings on different evaluations.

For an example of either of these theories in action, consider the analysis of

French schwa given in Chapter 1. The relevant constraints are *CNC, *ə, and MAX. Since no variant violates *CNC, it does not have a variable ranking in the hierarchy. The other two, however, do. This looks like (12).

(12) a. Tableau for French, *ə >> MAX

/ãvi də tə lə dəmãde/	*CNC	*ə	MAX
(←) a. ãvi də tə lə dəmãde		***!*	
(←) b. ãvi d_ tə lə dəmãde		***!	*
(←) c. ãvi də t_ lə dəmãde		***!	*
(←) d. ãvi də tə l_ dəmãde		***!	*
(←) e. ãvi də tə lə d_mãde		***!	*
→ f. ãvi d_ tə l_ dəmãde		**	**
→ g. ãvi də t_ lə d_mãde		**	**
h. ãvi də t_ l_ dəmãde	*!	**	**

b. Tableau for French, MAX >> *ə

/ãvi də tə lə dəmãde/	*CNC	MAX	*ə
→ a. ãvi də tə lə dəmãde			****
(←) b. ãvi d_ tə lə dəmãde		*!	***
(←) c. ãvi də t_ lə dəmãde		*!	***
(←) d. ãvi də tə l_ dəmãde		*!	***
(←) e. ãvi də tə lə d_mãde		*!	***
(←) f. ãvi d_ tə l_ dəmãde		*!*	**
(←) g. ãvi də t_ lə d_mãde		*!*	**
h. ãvi də t_ l_ dəmãde	*!	**	**

In (12a), *ə >> MAX. This ranking favors as much deletion as possible, which in the absence of *CNC would mean leave no [ə] in the optimal output form. As it is, candidates (f) and (g), which show the most deletion possible, are the tied winners, and none of the more-faithful candidates are preserved. In (12b), MAX >> *ə. This ranking favors NO deletion, so only candidate (a) wins – all of the candidates with any number of missing [ə] are killed by MAX.

As is apparent from (12), the Partial Orders theory permits only those candidates that BEST satisfy the highest ranking variable constraints to win. All of the candidates that

PARTIALLY satisfy MAX or *ə are thus collectively harmonically bounded (Samek-Lodovici & Prince 1999, 2005) by either (a), the fully-faithful candidate, or by (f) or (g), the most-[ə]-free candidates. Since, by definition, no harmonically bounded form can win, these candidates are all equally doomed. Consequently, these theories predict that variation cannot occur locally, at a single locus of variation: each choice of variant must be consistent throughout an entire candidate in order for that candidate to win. One of the chief motivations for seeking an alternative to the Partial Orders theory, like Markedness Suppression, is to find a solution for this problem.

Alternatively, Kaplan (2012, 2014) gives a possible solution WITHIN the Partial Orders theory, based on McCarthy's (1982) observations about this flapping in English, shown in (13).

(13) *repetitive*

a. repe[t^h]i[t^h]ive
b. repe[ɾ]i[t^h]ive
c. repe[ɾ]i[ɾ]ive
d. *repe[t^h]i[ɾ]ive

In English, /t/ can surface as either [ɾ] or [t^h] intervocalically. In a word with multiple instances of /t/, like *repetitive*, not all combinations of these two are possible: a word with all of one allophone or all of the other is fine, but [t^h] cannot precede [ɾ] (though [ɾ] can precede [t^h]). McCarthy attributes this to the prosodic structure of the word, summarized in (14): flapping is licensed in /peti/, where it is foot-internal (the domain labeled Σ in (14)), but can also apply to broader prosodic domains containing that foot, which in this case is /petitive/ (labeled Σ' in (14)).

(14) re[[peti $]_{\Sigma}$ tive $]_{\Sigma'}$

The Partial Orders theory can account for this by using constraints favoring flapping specific to each of the domains Σ and Σ' instead of a single constraint favoring flapping throughout the entire word.

Here, I will sketch out an application of this analysis to French, using *envie de te le demander* as an example. (For a full discussion, see Kaplan (2014).) This phrase has a rich prosodic structure: *de*, *te*, and *le* are all proclitics. An analysis in which constraints like *ə/X, where X is some aspect of this prosodic structure – a boundary, an entire domain, a dominance relationship, et cetera – replace the context-independent constraint *ə might be able to access all of the candidates that are harmonically bounded if *ə is used instead, as in (12).

A full analysis on the basis of this theory is beyond the scope of this chapter, but here I will show how just one of these context-sensitive versions of *ə changes the analysis for the better. In addition to the three constraints used in the analysis in (12), suppose we add *ə/PWD, which applies to any domain headed by a prosodic word. Then, the ranking in (15) is possible.

(15) Tableau for French, *ə/PWD >> MAX >> *ə

/ãvi də tə lə $[\text{dəmãde}]_{\text{PwD}}$/	*CNC	*ə/PWD	MAX	*ə
a. ãvi də tə lə $[\text{dəmãde}]_{\text{PwD}}$		*!		****
b. ãvi d_ tə lə $[\text{dəmãde}]_{\text{PwD}}$		*!	*	***
c. ãvi də t_ lə $[\text{dəmãde}]_{\text{PwD}}$		*!	*	***
d. ãvi də tə l_ $[\text{dəmãde}]_{\text{PwD}}$		*!	*	***
→ e. ãvi də tə lə $[\text{d\_mãde}]_{\text{PwD}}$			*	***
f. ãvi d_ tə l_ $[\text{dəmãde}]_{\text{PwD}}$		*!	**	**
g. ãvi də t_ lə $[\text{d\_mãde}]_{\text{PwD}}$			**!	**
h. ãvi də t_ l_ $[\text{dəmãde}]_{\text{PwD}}$	*!	*	**	**

In the tableau in (15), [dəmãde] is the only prosodic word, and thus, the domain of *ə/PWD is restricted only to the [ə] in [dəmãde]. Consequently, *ə/PWD eliminates only those candidates that preserve the [ə] in that position, indifferent to what those candidates do with [ə] at any other position. This leaves MAX to decide between the candidates, (e) and (g), that delete [ə] in that position. As MAX favors the minimal number of deleted segments, any candidate, like (g), that deletes more than the [ə] in [dəmãde] is eliminated, leaving only (e).

The addition of this constraint allows one of the previously harmonically bounded candidates in which only a single [ə], belonging to a prosodic word, has been deleted in order to satisfy *ə/PwD. The addition of more constraints of the *ə/X class can allow access to more of the harmonically bounded candidates. However, as Kaplan (2014) observes, it is not clear that this solution can be extended to ALL kinds of local variation, as it is dependent on the ability to categorize loci of variation into distinct domains. This remains a project fit for future research within the Partial Orders paradigm.

2.3 Rank-Ordered EVAL

Rank-Ordered EVAL (Coetzee 2004, 2006), in contrast to the Partial Orders theory, does not rely only on changes to CON. Instead, it also targets EVAL. The idea is that EVAL, instead of merely picking a winner, makes a ranking between ALL of the candidates created by GEN, and then one of the candidates is selected as the optimal output. Those candidates which are more "harmonic" (that is, higher ranked on this scale) are more likely to be chosen as winners. Accompanying this framework is a division between constraints that are or are not capable of eliminating candidates, usually separated with a cut-off line (represented with two vertical bars in (16) below). Candidates that violate constraints above the cut-off line will not be accessed so long as a candidate that does not violate a constraint above the cut-off line is available.

One potential analysis for French using Rank-Ordered EVAL is found in Kaplan (2011). Since *ə is what motivates variation in the first place, it must be ranked below the cut-off line. However, French does not permit deletion if it would create an illicit consonant cluster, so any constraints banning those – in our case, just *CNC – must be ranked above the cut-off line, as those can never be violated. MAX, on the other hand, has no phonotactic motivation and competes with *ə beneath the cut-off line for determining the well-formedness of outputs. The result of this ranking is given in (16).

(16) Tableau for French, ROE analysis

/ãvi də tə lə dəmãde/	*CNC		*ə	MAX
→ a. ãvi də tə lə dəmãde			****	
→ b. ãvi d_ tə lə dəmãde			***	*
→ c. ãvi də t_ lə dəmãde			***	*
→ d. ãvi də tə l_ dəmãde			***	*
→ e. ãvi də tə lə d_mãde			***	*
→ f. ãvi d_ tə l_ dəmãde			**	**
→ g. ãvi də t_ lə d_mãde			**	**
h. ãvi də t_ l_ dəmãde	*!		**	**

The only constraint that can eliminate candidates here is *CNC. Consequently, the only candidate that is eliminated is (h), the only one to violate *CNC. Each of the other candidates is a possible output form whose probability of success is influenced by the constraints below the cut-off line, which rank candidates with respect to how well they are satisfied. For example, candidate (g) in (16) is expected to be more frequently realized than candidate (a), as candidate (a) violates the higher ranking constraint *ə four times while (g) does so only twice. However, although the theory establishes relative frequencies for each output candidate, it does not make predictions about the ABSOLUTE frequency of any given output pattern as in Partial Orders, Stochastic OT, and Markedness Suppression theory. Coetzee argues that (a) it is not clear what frequencies are actually being modeled by these theories since individual speakers have different preferences for each output pattern and corpus averages may not correspond to any real rate of preference, and (b) variation is too heavily conditioned on nonphonological factors for such predictions to be valid as a basis of grammar (Coetzee 2004: 305-308).

As noted in the introduction, I find this argument convincing. Yet, as Markedness Suppression can make absolute predictions, I will provide those calculations wherever corpus frequencies are available to provide a comparison. The extent to which these comparisons matter, and the larger issue of the importance of absolute versus relative frequencies, is somewhat beyond the scope of this thesis.

2.4 Conclusion

The initial theories of variation, based on the possibility of variation in the constraint ranking itself, set the tone for variation in OT. Subsequent theories have all tried to improve on the predictions made by the Partial Orders theory. In particular, theories try to produce better predictions of variant frequencies (e.g. Stochastic OT) and/or more empirical coverage (e.g. Rank-Ordered EVAL), particularly in coming up with solutions to the harmonic bounding problem.

Markedness Suppression tries to improve on Partial Orders in two ways. Structurally, it is set up as an answer to the harmonic bounding problem, explicitly permitting harmonically bounded forms to surface by permitting suppression of violations to affect only a single locus of variation at a time. This also enables it to produce a broader range of variant frequencies than those permitted by Partial Orders, which is locked into only the probabilities allowed by the number of constraints subject to a variable ranking.

CHAPTER 3

COORDINATION BETWEEN LOCI: WARAO [p]~[b] VARIATION AND SHIMAKONDE VOWEL REDUCTION

3.1 Introduction

A cornerstone of the architecture of Markedness Suppression is the assumption that choices at each locus of variation are independent from choices made at any other locus. Thus, in the French schwa deletion pattern described in Chapter 1, the decision to delete one schwa is presumably indifferent to the decision to delete any other schwa. Some variation, however, appears to be the result of exactly this: loci work in tandem in such a way as to "lock in" certain variations at some loci of variation given choices that are made elsewhere, at other loci of variation.

Two languages have patterns that are particularly good examples of such coordinated variation. The first, Warao, permits free variation between [p] and [b], but does not permit combinations of the two in any given word (Osborn 1966). This has been used to argue in favor of a distinction between local and global variation (Kimper 2011a). The second, Shimakonde, permits optional reduction of pretonic midvowels to [a], so long as reduction begins at the left edge and does not skip any midvowels along the way (Liphola 2001).

This chapter argues that the analyses of each of these patterns under Markedness

Suppression is straightforward: while it is true that the constraint motivating variation, on its own, overgenerates, the introduction of a set of constraints characterizing the supplementary restrictions on variation present in each of these languages permits us to predict only the attested output forms. On this basis, I argue that this makes all patterns of coordinated variation formally equivalent to the avoidance of illicit structural configurations, just as French [ə] deletion was restricted to avoid phonotactic violations in Chapter 1. This predicts, further, that coordination in variable patterns is not arbitrary or random, but arises directly from other structural requirements in a language.

Section 3.2 will look at Warao, in which the supplementary pattern appears to be consonant harmony. Section 3.3 will handle Shimakonde, where ALIGN plays this role. Finally, Section 3.4 will generalize this approach and discuss its implications.

3.2 Warao

3.2.1 Analysis

Warao permits free variation between [p] and [b], but does not allow them to mix.

(17) /paro + parera/ ‘weak’

a. **p**aro**p**arera
b. **b**aro**b**arera
c. ***p**aro**b**arera
d. ***b**aro**p**arera

(Osborn 1966)

In addition to [p], Warao permits the voiceless obstruents [t k k^w s h]. Voiced obstruents are disallowed, with the exception of [b]. Consequently, *OBSVOICE must be high-ranking. We can then motivate the choice of [b] over [p] by ranking Θ *P over *OBSVOICE, where the optionality of *P (denoted by Θ) is used to motivate variation.

As predicted, this ranking overgenerates, shown in (18).

(18) Tableau for /paro + parera/ ‘weak’ (overgeneration)

/paro + parera/	Θ *P	*ObsVoice
(→) a. paroparera	**	
(→) b. barobarera		**
💣 c. parobarera	o	*
💣 d. baroparera	o	*

The problem is that there is no constraint in (18) to rule out (c) and (d), which surface when either of their violations of Θ *P are suppressed due to having one fewer voiced obstruent than (b). The solution is to narrow down the candidates that Θ *P can decide among by categorically eliminating those like (c) and (d) that mix [p] and [b]. Since we want to enforce the same specification for [voice] at the loci of variation in (c) and (d), this appears to be a fairly straightforward case of consonant harmony.

I will account for consonant harmony according to Hanson (2001), who uses an approach based on Agreement by Correspondence (Walker 2000a,b, 2001a) that he refers to as "consonant-consonant correspondence" (CC-correspondence) to achieve this. The idea behind this theory, motivated by psycholinguistic research on speech planning, is that similar consonants – those that share particular features – stand in correspondence with one another in output strings. They can then be required by faithfulness constraints to have the same specification for certain features, just as with input-output or output-output correspondence.

Formally, correspondence is enforced by the constraint in (19).

(19) Corr[X] ≡
Consonants in the output with the same specification for [X] must be in correspondence.

The feature that [p] and [b] share but Warao's other obstruents do not is [+labial], so the correspondence constraint I will use here is (20).

(20) CORR[LABIAL] ≡
[+labial,+obstruent] consonants in the output must be in correspondence.

In order to handle similar patterns of voicing harmony in Ngbaka (a Gbaya language spoken in the Congo) and Chara (a Semitic language spoken in Ethiopia), Rose and Walker (2001) propose the constraint in (21), which requires identical specification for voicing among corresponding consonants.

(21) IDENT[VOICE]-CC ≡
Consonants in correspondence in the output must have the same specification for [voice].

Θ *P, IDENT[VOICE]-CC, and CORR[LABIAL] have no necessary ranking between them, as they don't interface with one another. All three constraints must outrank *OBSVOICE: Θ *P because otherwise it would have no effect, and the other two in order to prevent circumvention of the agreement-by-correspondence mechanism since *OBSVOICE favors the illicit mixed forms over those with all [b]. This ranking is shown in (22).

(22) Tableau for /paro + parera/ ‘weak’ (no suppression)

/paro + parera/	Θ *P	CORR[LABIAL]	IDENT[VOICE]-CC	*OBSVOICE
(→) a. $p_i$aro$p_i$arera	**			
→ b. $b_i$aro$b_i$arera				**
c. $p_i$aro$b_i$arera	*		*!	*
d. $b_i$aro$p_i$arera	*		*!	*
e. $b_i$aro$p_j$arera	*	*!		*

Proceeding from the penultimate candidate to the first, (c) and (d) are eliminated by IDENT[VOICE]-CC because their obstruents have not harmonized for voice. Both (a) and (b) are potential outputs here. If no violations of Θ *P are suppressed for (a), (b) wins, as (a) would be eliminated by that constraint. However, when both of (a)'s violations of Θ *P are suppressed, that is the winning candidate, as it incurs no violations of *OBSVOI while (b) incurs two, as shown in (23).

(23) Tableau for /paro + parera/ ‘weak’ ((a) is the winner)

/paro + parera/	Θ *P	CORR[LABIAL]	IDENT[VOICE]-CC	*OBSVOICE
→ a. $p_i$arop$_i$arera	oo			
b. $b_i$arob$_i$arera				*!*
c. $p_i$arob$_i$arera	*!		*	*
d. $b_i$arop$_i$arera	*!		*	*
e. $b_i$arop$_j$arera	*!	*		*

The last candidate, (e), demonstrates why CORR[LABIAL] must outrank *OBSVOICE. If that were not the case**,** it would be possible to have an unattested form like (e) surface: with its violation of Θ *P suppressed and no more than two violations of (a)'s suppressed, only (e) and (b) would survive up to *OBSVOICE, which would then eliminate (b) and select (e) as optimal.

Similarly, in the rankings in (24) on the following page, it would be possible for (c) and/or (d) to win if their violations of Θ *P are suppressed because IDENT[VOICE]-CC is ranked below *OBSVOICE.

Finally, this ranking does not permit any other voiced obstruents to sneak into the output, which we would not expect anyway, as the constraint motivating the appearance of [b], Θ *P, is specific to [p].

(25) Tableau for /koyakitane/ 'to tie up'

/koyakitane/	Θ *P	CORR[LABIAL]	IDENT[VOICE]-CC	*OBSVOICE
→ a. $k_i$oyak$_i$it$_j$ane				
b. $g_i$oyag$_i$it$_j$ane				*!
c. $k_i$oyag$_i$it$_j$ane			*!	*
d. $g_i$oyak$_i$it$_j$ane			*!	*

(24) a. Tableau for /paro + parera/ ‘weak’ (*OBSVOICE >> CORR[LABIAL], (e) is winner)

/paro + parera/	Θ *P	*OBSVOICE	CORR[LABIAL]	IDENT[VOICE]-CC
a. $p_i$arop$_i$arera	*!*			
b. $b_i$arob$_i$arera		**!		
c. $p_i$arob$_i$arera	*!	*		*
d. $b_i$arop$_i$arera	*!	*		*
💣 e. $b_i$arop$_j$arera	o	*	*	

b. Tableau for /paro + parera/ ‘weak’ (*OBSVOI >> ID[VOICE]-CC, (c) & (d) winners)

/paro + parera/	Θ *P	*OBSVOICE	CORR[LABIAL]	IDENT[VOICE]-CC
a. $p_i$arop$_i$arera	*!*			
b. $b_i$arob$_i$arera		**!		
c. $p_i$arob$_i$arera	o	*		*
d. $b_i$arop$_i$arera	o	*		*
e. $b_i$arop$_j$arera	*!	*	*	

To conclude, although the constraint used to create variation, Θ *P, could only handle local variation on its own, thus overgenerating, the combination of this constraint and the consonant harmony constraints is capable of creating global variation by eliminating the unattested intermediate combinations of [p] and [b] permitted by Θ *P.

3.2.2 Discussion

Markedness Suppression differs from variable ranking theories of variation in that output probabilities are not directly tied to the number of possible combinations of violations, but rather to the probability with which those violations are suppressed. In this case, [p] will occur with probability p^n, while [b] occurs with probability $1 – p^n$, where p is the rate at which violations of Θ *P are suppressed and n is the number of violations incurred by the [p] variant (i.e. the number of /p/ in the word). Since these two probabilities sum to 1, we can simply choose p^n to approximate the rate at which the [p] variants are used, and we will have completely and accurately described the distribution of the variants.

One objection to this analysis may be the unavoidable prediction that [p] increases in rarity with n since the probability of [p], p^n, is monotonically decreasing in n. Though Osborn (1966) does not report a difference in frequency proportional to the number of obstruents, he also does not report any positive account of output frequencies; the only observation he makes with respect to general frequencies between the two variants is that [b] is more frequent, which is compatible with our analysis. There are, however, no clear data available to support or contradict this prediction.[6]

Interestingly, Osborn (1966) suggests that different classes of words (e.g. frequent words, words with Spanish origins) show favoritism for one variant over another. In the absence of clear data on the frequencies involved, it is not clear what to make of this; the solution may be to permit lexical selection of [p] or [b], suggesting an underlying unspecified consonant or perhaps restriction of constraints to particular lexical strata, per Ito and Mester (1999a). I leave this to future research on Warao, pending the availability of more data on these frequencies.

A second objection may be that this analysis is arbitrary: as Warao only has a single voiced obstruent, there is no way to falsify claims of voicing harmony and no way to judge if it is what is actually going on in this language. From the perspective of theory, however, this is not a problem. Voicing harmony is not IN GENERAL unattested (Rose & Walker 2001), and the Warao pattern is predicted by the intersection of our two theories: since Markedness Suppression predicts the possibility of variation between [p] and [b] – represented by Θ *P in our analysis – and consonant harmony predicts languages where some subset of consonants agree in voicing, the existence of both theories predicts a language like Warao, where the choice of [p] or [b] is consistent for an entire word.

Furthermore, Warao also demonstrates nasal harmony (Peng 2000), suggesting the class of constraints involved in harmony is already active in the language. The use of

[6] In the event that such data were produced, and those data were incompatible with the prediction with respect to the number of instances of [p], one solution would be to use a constraint that categorically bans [p], that is, a constraint that assigns only one violation regardless of the number of instances of [p] in a candidate.

consonant harmony to achieve the required analysis, though seemingly artificial, is not in any way contradicted by theory or data, and it does not appear to make any faulty predictions. Ultimately, the globality of the variation is in itself evidence of consonant harmony.

3.3 Shimakonde

3.3.1 Analysis

Shimakonde has optional pretonic (effectively prepenultimate) vowel reduction from midvowels to [a] that spreads right from the left edge of a word without skipping, a pattern Liphola (2001) calls "contiguity of reduction" (CORE).

(26) kú-pélévéléˈlééla ‘to not reach a full size for’

a. kú-pélévéléˈlééla
b. kú-p**á**lévéléˈlééla
c. kú-p**á**l**á**véléˈlééla
d. kú-p**á**l**á**v**á**léˈlééla
e. kú-p**á**l**á**v**á**l**á**ˈlééla
f. *kú-pél**á**v**á**l**á**ˈlééla
g. *kú-pélév**á**l**á**ˈlééla
h. *kú-pélévél**á**ˈlééla
i. *kú-p**á**lév**á**léˈlééla

(Liphola 2001)

In order to account for reduction, I will use Θ *UNSTRESSED/MIDV (Crosswhite 2001), which bans midvowels from all but the penultimate syllable. Since it is suppressible, this will allow for variation between faithful preservation of midvowels and reduction.[7]

Naturally, Θ *UNSTRESSED/MIDV still overgenerates on its own, shown in (27).

---

[7] This technically permits reduction by lowering or by raising. In order to prevent raising, MAX[-HIGH] can outrank Θ *UNSTRESSED/MIDV (Crosswhite 2001). This does not impact the interaction between CORE and Θ *UNSTRESSED/MIDV, so it has been omitted from tableaux in this section.

(27) Tableau for /-pélévéléˈlééla/ ‘to not reach a full size for’ (overgeneration)

/-pélévéléˈlééla/	Θ *UNSTRESSED/MIDV	IDENT[HEIGHT]
(→) a. -pélévéléˈlééla	******	
(→) b. -p**á**lévéléˈlééla	*****	*
(→) c. -p**á**l**á**véléˈlééla	****	**
→ d. -p**á**l**á**v**á**l**á**ˈlééla		****
(💣) e. -pélévél**á**ˈlééla	*****	*
(💣) f. -p**á**lév**á**léˈlééla	****	**

The problem, of course, is that there is no constraint(s) to enforce the parameters of CORE, which leaves (e) and (f), which follow the variable reduction part of the pattern just fine, as possible optimal candidates. In particular, (b) and (e) and (c) and (f) are equally acceptable under this constraint hierarchy. The solution is once again to narrow down the set of optimal outputs by eliminating candidates that do not follow CORE.

CORE presents candidates with two requirements: (i) that reduction begins on the left edge and (ii) that reduction proceeds rightward in an unbroken sequence. I will deal with these two requirements separately, as there is no reasonable constraint that enforces both in this case.

In order to account for the alignment of reduction, I will use the local conjunction (Smolensky 1993) of ALIGNL and IDENT given below.

(28) ALIGNL[+LOW] $\&_{C}$ IDENT[HEIGHT] ≡
Assign one violation for any [low] feature in the output that is not present in the input for each syllable between the syllable containing it and the left edge (of the prosodic word).

This constraint requires some justification. On its own, ALIGNL[+LOW] together with a positional faithfulness constraint to prevent tonic and posttonic syllables from raising in order to avoid violations of this constraint – I will use IDENT[HEIGHT]-FT for this purpose – works only for words with no underlying pretonic [a] somewhere between the first and the tonic syllable. For words that DO have an underlying [a], it fails, per (29).

(29) Tableau for /lekaˈniila/ 'leave for each other'

/lekaˈniila/	ʘ *UNSTRESSED/MIDV	IDENT[HEIGHT]-FT	ALIGNL[+LOW]	IDENT[HEIGHT]
← a. l**e**k**a**ˈniila	*		***,*!	
b. l**a**k**e**ˈniila	*		***	**!
💣 c. l**e**k**e**ˈniila	oo		***	*
→ d. l**a**k**a**ˈniila | / [+low]			***	*
e. l**e**k**a**ˈniil**e**	***	*!		***

The optimal candidates in (29) should be (a) and (d), since underlying /e/ can reduce to [a], but underlying /a/ cannot raise to [e]. Candidate (e) is killed by IDENT[HEIGHT]-FT, as it raises posttonically to avoid violations of ALIGNL[+LOW]. (For future tableaux, I will not consider candidates of this kind, as IDENT[HEIGHT]-FT is sufficient to handle them.) Candidate (d), appropriately, is optimal and receives no violations from ALIGNL (except for three from the final vowel) because the [a] in the second syllable shares its [+low] feature with the vowel in the first syllable. In contrast, (a) is wrongfully eliminated by ALIGNL because of the faithfully preserved [+low] feature in the second syllable. Additionally, candidate (c) is erroneously marked as optimal because it raises the first [a], trading in the extra violation of ALIGNL that killed (a) for a suppressible violation of ʘ *UNSTRESSED/MIDV that allows it to tie with (d).

This means there must be, at minimum, an IDENT constraint like IDENT[+LOW] outranking ALIGNL in order to prevent ALIGNL from penalizing the underlying [a]. In particular, since only the height of underlyingly low vowels is of interest, a directional faithfulness constraint (Pater 1996) like (30) is appropriate.

(30) IDENTI→O[+LOW] ≡
Assign one violation for each output correspondent of a [+LOW] segment in the input that is not [+LOW] in the output.

Since this is a directional constraint, it will not penalize any midvowels that reduce to [a] in the output; it will only assign violations if underlyingly low vowels are raised, as desired. Since this performs the same service as IDENT[HEIGHT]-FT, it can take that constraint's place.

Though this succeeds in correctly ruling out candidates like (b) and (c), it encourages full reduction as in (d) so that the [+low] feature appears in the first syllable, ruling out the (attested) fully faithful form (a). This is shown in tableau (31).

(31) Tableau for /lekaˈniila/ 'leave for each other'

/lekaˈniila/	Θ *UNSTRESSED/MIDV	IDENTI→O[+LOW]	ALIGNL[+LOW]	IDENT[HEIGHT]
← a. l**e**k**a**ˈniila	*		***,*!	
b. l**a**k**e**ˈniila	*	*!	***	**!
c. l**e**k**e**ˈniila	**	*!	***	*
→ d. l**a**k**a**ˈniila | / [+low]			***	*

There are two ways to get the fully faithful form to surface: either change the alignment constraint to prevent it from penalizing the underlying [a] or add an additional suppressible constraint to further penalize (d), but not (a). Any constraint of this kind would have to outrank ALIGNL[+LOW] and prefer nonreduction to reduction without interfering with ALIGNL[+LOW] in longer words, which seems like a tall order.[8]

The alternative is to modify the ALIGNL[+LOW] constraint to make its requirements less strict. Conjoining ALIGNL with an IDENT constraint prevents it from assigning an additional violation to the unreduced form, preventing it from killing off (a) and successfully generating the attested set of optimal candidates.

[8] One such constraint could be CRISPEDGE (Ito & Mester 1999b, Walker, 2001b), which penalizes the spreading of features. This would work for words with two pretonic vowels, where skipping one is not possible, but it would not work with words with more pretonic vowels: in order to prevent ALIGNL[+LOW] from killing candidates with offset underlying low vowels, it would have to outrank it, but this encourages avoiding violations of CRISPEDGE by skipping vowels in longer words instead of feature sharing, contrary to CORE.

(32) Tableau for /lekaˈniila/ 'leave for each other'

/lekaˈniila/	Θ *UNSTRESSED/MIDV	IDENTI→O[+LOW]	ALIGNL[+LOW] &C IDENT[HEIGHT]
(→) a. l**e**k**a**ˈniila	*		
b. l**a**k**e**ˈniila	*	*!	
c. l**e**k**e**ˈniila	**	*!	
→ d. l**a**k**a**ˈniila | / [+low]			

This also correctly accounts for words in which only a single vowel could reduce.

As discussed above, when multiple midvowels precede the tonic syllable, CORE also requires that they reduce in an unbroken sequence. Formally, the conjoined alignment constraint above already favors sharing of the [+low] feature that originates on the left edge, as additional copies of the [+low] feature further to the right will incur additional violations of ALIGNL. However, skipping a syllable does not necessitate NEW copies of [+low]. Walker (2011) discusses, in the context of licensing patterns, long-distance feature association by correspondence, as shown in (33).

(33) –p**á**lév**á**léˈlééla
| |
[+low]$_i$ [+low]$_i$

The structure in (33) is a "feature chain" in which what is essentially a single [+low] feature appears in two separate instances in nonadjacent syllables linked not by sharing, but by a correspondence relationship between instances of the feature. This structure would not receive any violations of the conjoined alignment constraint because the feature appears in the first syllable, satisfying the left-edge requirement.

To penalize this structure, Walker (2011) uses a constraint of the form (34), which simply prevents the correspondence relationship in (33).

(34) *DUPLICATE[+LOW] (= *DUPLICATE) ≡
Copies of [+low] cannot stand in correspondence.

The effect of this constraint is shown in (35).

(35) Tableau for /-pélévéléˈléela/ ‘to not reach a full size for’

/-pélévéléˈléela/	Θ *UNSTRESSED /MIDV	*DUPLICATE	IDENTI→O[+LOW]	ALIGNL[+LOW] &$_C$ IDENT[HEIGHT]
(→) a. –pélévéléˈléela	****			
(→) b. –p**á**lévéléˈléela	***			
(→) c. –p**á**l**á**véléˈléela	**			
(→) d. -p**á**l**á**v**á**l**á**ˈléela				
e. -pélévél**á**ˈléela	***			*!**
f. -p**á**$_i$lév**á**$_i$léˈléela	**	*!		
g. -p**á**$_i$lév**á**$_j$léˈléela	**			*!*

Here, (e) and (g) are eliminated by ALIGNL[LOW] &$_C$ IDENT[HEIGHT]. For (g), the two pretonic [low] in the output are NOT part of the same feature chain, as indicated by the mismatched subscripts, but are actually part of two different chains. The conjunction constraint thus assigns two violations to the second [low]. Candidate (f) reflects the opposite relationship between the two [low] features; these are in the same chain. Accordingly, the conjunction constraint assigns no violations, but the candidate is eliminated by *DUPLICATE. The other four candidates are all optimal depending on the way violations of Θ *UNSTRESSED/MIDV are or are not suppressed.

To conclude, the analysis given above is capable of accounting for both the variability of reduction as well as the structural restrictions of CORE. In the next section, I will show that this analysis is also compatible with the opaque interactions between reduction and vowel harmony and vowel coalescence in Shimakonde.

3.3.2 Reduction with harmony and coalescence

Reduction interacts opaquely with two other processes in the language (Ettlinger 2009): vowel harmony (where reduction results in overapplication of harmony) and vowel coalescence (where reduction results in underapplication of coalescence).

Consequently, any analysis of reduction must also be compatible with an account of these two processes. In this section, I will show that this is the case for the analysis given above. First, I will discuss vowel harmony.

Vowel harmony causes [-low] segments to harmonize their height features (in particular, high vowels lower to midvowels when harmonizing with midvowels). Harmony is opaque when reduction results in outputs where mid vowels have been reduced, but the high vowels still harmonize as if the midvowels have not been. This is shown in (36).

(36) Vowel harmony (opaque interaction)
a. /va-ndá-t**o**t-**í**l-a/ → vandat**á**ˈt**éé**la, *not* *vandát**á**ˈt**íí**la 'they will sow for'
b. /va-nda-ʃ**é**m-**í**la/ → vandaʃ**á**ˈm**éé**la, *not* *vandaʃ**á**ˈm**íí**la 'they will call for'
Liphola (2001)

In (36a), for example, the midvowel /o/ in the input reduces to [a] in the output. Since high vowels harmonize with mid, not low, vowels, the transparent outcome would leave /i/ surfacing faithfully as [ii]. This is not what happens: instead, it lowers to [ee] as if harmonizing with the midvowel in the input.

Ettlinger (2009) uses "diagonal correspondence" to account for this kind of opaque harmony. The basic intuition behind this is that correspondence occurs not only between segments in the output, but also those in the input. Since this theory is used to implement harmony, it is specified to nonlow vowels. Thus, in (36a), the /o/ and /i/ in the input as well as the [ee] in the output would all stand in correspondence with one another; in (35b), the /e/ and /i/ in the input and the [ee] in the output would also stand in correspondence. The constraint that achieves this is (37).

(37) CORRD[-LOW] ≡
All [-low] segments in both the input and output are in correspondence with one another.

As with agreement by correspondence in Warao, we can use IDENT[HEIGHT]-CC to force corresponding vowels to harmonize for height. The cover constraint Θ REDUCE, which I will use as a stand-in for the analysis in Section 3.1, can then produce the opaque interaction.

(38) Tableau for /va-ndá-tot-íl-a/ 'they will sow for'

/va-ndá-t**o$_i$**t-**í$_i$**l-a/	Θ REDUCE	CORRD[-LOW]	IDENT[HEIGHT]-CC	IDENT[HEIGHT]
a. vandát**o$_i$**ˈt**íí$_i$**la	*		*!	
b. vandát**a**ˈt**íí$_i$**la			*!	*
→ c. vandát**o$_i$**ˈt**ee$_i$**la	*			*
→ d. vandát**a**ˈt**ee$_i$**la				**
e. vandát**a**ˈt**íí$_i$**la		*!		

The only thing that differentiates the optimal candidates (c) and (d) in (38) is whether or not reduction occurs; both follow vowel harmony correctly. The other three candidates are eliminated because of the vowel harmony constraints: (e) because it does not correctly reflect correspondence relations and (a) and (b) because corresponding segments do not have matching height features.

The other process we need to consider is vowel coalescence, which occurs as a means of hiatus resolution. Reduction is transparent with respect to the coalescence in (a) (vacuously, as reduction does not target low vowels) and (b), but midvowels resulting from the coalescence in (c) do not reduce. This is shown in (39).

(39) Vowel coalescence

a. Low + Low → Low
/va-nd**a-á**kat-a/ → vanˈd**aá**kata 'we will move'

b. Low + Mid → Mid
/tu-nd**a-é**k-a/ → tunˈd**eé**ka 'we will laugh'

c. Low + High → Mid
/tu-nd**a-í**m-a/ → tunˈd**eé**ma 'we will deny'

Liphola (2001)

(40) Reduction + Vowel coalescence

a. Transparent

/va-nd**a-é**k-an-a/ → vand**eé**ˈkáána *or* vand**aá**ˈkáána 'they will laugh (at) each other'

/va-nd**a-é**p-an-a/ → vand**eé**ˈpaána *or* vand**aá**ˈpaána 'they will harvest each other'

b. Opaque

/va-nd**a-í**m-an-a/ → vand**eé**ˈmaána, *but* *vand**aá**ˈmaána 'they will deny each other'

/va-nd**a-i**tík-a/ → vand**eé**ˈtiíka, *but* *vand**aá**ˈtiíka 'they will respond'

Ettlinger (2009)

We can account for hiatus with $*V_iV_j$, which penalizes adjacent vowels of different quality. This must be outranked by MAX to rule out deletion as a repair strategy. In both the transparent and opaque cases, if there is a high or mid vowel present, that quality is preserved over the low vowel by coalescence, which can be represented by MAX[-LOW]. However, in the transparent cases, reduction still occurs, so the reduction constraints – represented here by Θ REDUCE once again – must outrank it. These constraints suffice for the transparent case:

(41) Tableau for /va-nda-ék-an-a/ 'they will laugh (at) each other'

/va-nd**a-é**k-an-a/	MAX	$*V_iV_j$	Θ REDUCE	MAX[-LOW]
→ a. vand**aá**ˈkáána				*
(→) b. vand**eé**ˈkáána			*	
c. vand**aé**ˈkáána		*!	*	
d. vand**a**ˈkáána	*!			

Meanwhile, high vowels always lower to mid, suggesting that the [-high] feature of the low vowel is preserved (a fact that is not reflected on the surface of the transparent cases). This can be accounted for with MAX[-HIGH], which needs to outrank MAX[-LOW].

High vowels, however, cannot lower further to low vowels; hence the opaque interaction with reduction. This can be analyzed as a ban on chain shifting from a high to a mid to a low vowel, which in OT can be accounted for with local constraint conjunction (Kirchner 1996). In particular, we want to prevent multiple features of underlying high vowels from changing at once, so the relevant constraint here is IDENT[+HIGH] $\&_C$ IDENT[-LOW], which must outrank the reduction constraints.

Together, these additional constraints account for the opacity here, as shown in (42).

(42) Tableau for /va-nda-itík-a/ 'they will respond'

	/va-nd**a-i**tík-a/	IDENT[+HIGH] $\&_C$ IDENT[-LOW]	MAX[-HIGH]	$*V_iV_j$	Θ REDUCE	MAX[-LOW]
	a. vand**aá**ˈtiíka	*!				*
→	b. vand**eé**ˈtiíka				**	
	c. vand**aé**ˈtiíka			*!	*	
	d. vand**ií**ˈtiíka		*!			

Thus, the analysis of reduction in Section 3.2 is compatible with both the vowel harmony and vowel reduction.

3.3.3 Discussion

Leaving opacity, this analysis incurs the same potential empirical problem encountered with Warao, namely, that unreduced forms are predicted to become rarer with the number of possible loci for reduction (the number of pretonic midvowels). This is due to the fact that, for any unreduced form to surface optimally, p^n violations of Θ *UNSTRESSED/MIDV must be suppressed, where n is the number of midvowels in the input, and this number decreases exponentially as n goes up. Whether or not this is actually a problem for the analysis depends on whether or not this is an accurate description of the data. In the absence of those data, this question must be left unanswered.

3.4 Conclusion

I have shown that the variation in the Warao data can be accounted for using Markedness Suppression, while the failure to mix [p] and [b] can be accounted for by consonant harmony. I have also shown that Shimakonde's spreading can be accounted for by aligning any possible reductions to the left and optionally reducing from there.

These two cases provide a paradigm for analyzing coordinated variation through Markedness Suppression. Optional constraints, on their own, will overgenerate and permit candidates to ignore coordination between loci. That coordination can then be enforced by a separate set of constraints to contract the set of optimal candidates to what is actually attested. So long as coordination can be separated into variation between a set of features and an additional restriction on where those features can occur, this method will suffice.

This means that coordination is no different from simply requiring that variation does not create illicit structural configurations. In other words, the patterns of Warao and Shimakonde, although they are the results of the application of a particular process (voicing harmony and CORE), are formally identical to French [ə] deletion. To see this, consider this tableau, the first shown in Chapter 1.

(43) Tableau for *envie de te le demander* 'feeling like asking you for it'

/ãvi də tə lə dəmãde/	*CNC	Θ *ə	MAX
(→) a. ãvi də tə lə dəmãde		****	
(→) b. ãvi d_ tə lə dəmãde		***	*
(→) c. ãvi də t_ lə dəmãde		***	*
(→) d. ãvi də tə l_ dəmãde		***	*
(→) e. ãvi də tə lə d_mãde		***	*
→ f. ãvi d_ tə l_ dəmãde		**	**
→ g. ãvi də t_ lə d_mãde		**	**
h. ãvi də t_ l_ dəmãde	*!	**	**

Notice that, in each of these candidates, there are points at which deletion of one [ə] is bled by the deletion of another [ə] elsewhere. This is best illustrated in (h), which is illicit because either the [ə] in [də] prevents the deletion of the [ə] in [tə], or vice-versa. Although there is no underlying pattern to the variation here beyond avoidance of certain consonant clusters, as there is in both Warao's "global variation" and Shimakonde's CORE, the loci of variation in French are ultimately just as "coordinated" as they are in Warao and Shimakonde. The paradigm for analyzing those two languages, described above, is identical to the paradigm for analyzing French.

Ultimately, "coordination" can be analyzed as the combination of two separate patterns: variation between some set of features plus the supplementary requirements on variation. Each of these patterns can be captured with a different set of constraints so that ultimately the variable markedness constraint(s) only decide between otherwise well-formed output. The only way in which this differs from any other analysis is that the supplementary pattern is more complex than a simple set of markedness constraints. This implies that there is no need to have a concept of "coordinated" or "global" variation, at all. The variation component of the pattern is always the same.

This implies, of course, that coordinated variation is never coordinated in arbitrary ways: for example, there should not be a pattern in the style of Warao where [p] and [b] cannot be mixed with, say, two or fewer instances of underlying /p/, but once there are three or more, one can mix and match laryngeal features to his or her heart's content. Since the mechanism for variation under Markedness Suppression is the random and independent suppression of violations, coordination is predicted to be a result of structural restrictions/phonotactics that can be accounted for by the constraint hierarchy, so we would not expect any patterns that cannot be cleanly captured in this way.

With respect to Markedness Suppression, our concern going forward is not whether or not it can be used to account for these patterns, as it clearly can. Instead, the question that remains is whether the empirical predictions we have made with respect to the frequencies (in particular, the relative frequencies) of particular forms holds up. Since this requires data on the frequencies of each form that is currently unavailable, this is an issue for future research to tackle.

CHAPTER 4

INTERACTION BETWEEN ENVIRONMENTS: HUNGARIAN VOWEL HARMONY AND ENGLISH T/D DELETION

4.1 Introduction

Variation is often conditioned by environment; that is, certain phonological environments are more or less permissive for one or more variants. Kaplan's (2011) analysis of French, for example, observes asymmetries in the frequency of deletion of [ə] arising from the kind of consonant clusters that deletion would create as well as their position in the word or phonological phrase.

This kind of conditioning does not NECESSARILY pose a problem for Markedness Suppression, as it does not in itself violate any of the theory's basic assumptions. It is nevertheless a potential obstacle for Markedness Suppression because of the kinds of analysis that must be adopted to account for these asymmetries, i.e. those with multiple (suppressible) constraints. This obstacle can be overcome simply by producing a Markedness Suppression analysis of these patterns, but in addition to that, I intend to show that it is GENERALLY possible to come up with such an analysis.

Markedness Suppression makes two fairly strong predictions about the interaction between a suppressible constraint and the other (suppressible) constraints in the ranking: (i) increasing the number of suppressible constraints – and, by extension, the number of suppressible violations – will affect the probability of a candidate's success by a factor of p^n, n being the number of additional violations that need to be suppressed (or not), and (ii) that suppressible constraints should always be "variable" with respect to every constraint they dominate because any combination of candidates that would ordinarily be eliminated by suppressible constraints can survive given that the relevant killing violations are suppressed, so it is impossible for a suppressible constraint to dominate another constraint in the usual sense. Instead, the ranking of a suppressible constraint over a fixed constraint merely guarantees no candidates are eliminated by the fixed constraint before they encounter the suppressible one.[9]

In systems where there is a potentially large amount of interdependence among environments, and consequently the constraints that represent them, this could add up to being a problem. Two systems that may fit this description are vowel harmony in Hungarian and word-final t/d deletion in some dialects of English. In Hungarian, front unrounded vowels, typically neutral in the harmonic system, contribute to variable harmony when following back vowels in a stem, but not in any other configuration. Constraints favoring or disfavoring backness harmony may therefore compete with one another in a manner sensitive to the dominance relation among them. In English, t/d deletion is variably conditioned by the following segment or pause, and the degree to which deletion is encouraged depends on the dialect. Previous analyses of this phenomenon undertaken with the variable rankings approach (e.g. Coetzee 2004,

---

[9] For example, consider a ranking of the form

$$\Theta\ X >> Z,\ Y,$$

where Θ X is a markedness constraint and Y and Z are any constraints. It is impossible to guarantee that Θ X will eliminate any candidates before Z and Y decide between them; Θ X does not reliably behave as if it dominates them.

Kiparsky 1993) have modeled this using at least three constraints and their factorial typology, which Markedness Suppression cannot immediately replicate.

This chapter will show that neither of these systems actually prove problematic for Markedness Suppression per the dominance problem (ii), described above, provided an accommodating analysis. On the contrary, the analyses proposed here are just as, if not more, capable as analyses proposed under competing theories of phonological variation. In order to make this argument, I will propose analyses for Hungarian and English in which dominance is NOT necessary and, consequently, Markedness Suppression succeeds.

The generalization I intend to suggest is that, while analyses based on domination use the structural relationship of the constraints to determine which environment is subject to which restrictions (and at what rate) since Markedness Suppression internalizes variability within each suppressible constraint, this effect can also be achieved by using constraints that are constrained to specific environments. Given constraints of this kind, stringency (de Lacy 2003) can be used to produce a relationship between them, which, due to the mechanics of Markedness Suppression, corresponds to the attested differentiation of frequencies in each environment. Since stringency relies on the existence of a natural hierarchy between each environment, this predicts that any differentiation of this kind is caused by and/or follows a natural hierarchy.

The conclusion with respect to the frequency problem (i) is less favorable since although Markedness Suppression can generate the correct output patterns, it does not have enough flexibility to model frequencies correctly in Hungarian. On the other hand, both the absolute and especially the relative frequency predictions made for English are fairly close to those attested in a variety of dialects.

The next section will provide an analysis of Hungarian vowel harmony. Section 4.3 will then take a look at English t/d deletion. Finally, Section 4.4 generalizes the approaches taken toward these two languages and discusses its implications.

4.2 Hungarian vowel harmony

4.2.1 Introduction

The account of Hungarian vowel backness harmony presented here is based primarily on Hayes and Londe (2006), which gives both a formal analysis (using stochastic OT to account for variation) as well as an empirical estimate of the frequency of each variant, attained through a study of native speaker responses to wug tests.

Hungarian vowels fall into two major categories: front and back harmonic vowels on the one hand and harmonically neutral vowels on the other. The front harmonic vowels are all [-back, +round]: the high vowels [y yː] and the mid vowels [ø øː]. The back harmonic vowels are all, correspondingly, [+back], and most are [+round]: high vowels [u uː], mid vowels [o ɔ oː], and the low vowel [aː], which is also unrounded. The neutral vowels are all [-back, -round]: the "low" vowel [ɛ], which is labeled low due to the height effect (see below), the mid vowel [eː], and the high vowels [i iː].

Stem-internal vowels are not subject to harmony. Instead, suffix vowels harmonize with the last vowel in the stem, as shown in (44). In the examples below, and all those that follow, harmonic front vowels will be represented with F, harmonic back vowels B, and neutral vowels N. Hence, a word like (44ai) [ɔblɔk-nɔk] *ablaknak* can be schematized BB-B, as the stem contains two back vowels and the suffix vowel is also

[+back].

The combinations of stem vowels listed in (44) take completely predictable suffix vowels.

(44) Predictable harmony

a. Transparent [+back] harmony

i. BB	[ɔblɔk-nɔk]	*ablaknak*	'window-DAT'
ii. NB	[biːroː-nɔk]	*bírónak*	'judge-DAT'
iii. FB	[glykoːz-nɔk]	*glükóznak*	'glucose-DAT'

b. Transparent [-back] harmony

i. F	[yʃt-nɛk]	*üstnek*	'cauldron-DAT'
ii. BF	[ʃoføːr-nɛk]	*sofőrnek*	'chauffer-DAT'

c. [-back] harmony with FN(N) stems

i. FN	[fyːsɛr-nɛk]	*füszernek*	'spice-DAT'
ii. FNN	[øːrizɛt-nɛk]	*őrizetnek*	'custody-DAT'

d. Lexical specification with all-N stems

i. N-F	[kɛrt-nɛk]	*kertnek*	'garden-DAT'
ii. N-B	[hiːd-nɔk]	*hídnak*	'bridge-DAT'

Hayes & Londe (2006)

In (44a,b), the last vowel in the stem belongs to one of the classes of harmonic vowels, F or B. In those cases, the rest of the vowels in the stem are irrelevant: the suffix vowel simply harmonizes with the final vowel of the stem. In (44c), though the last vowel in the stem is harmonically neutral, the suffix reliably harmonizes with the back vowel. For (44d), there is no word-level variation between choice of suffix vowels: a given word will reliably choose a particular vowel. Following Hayes and Londe (2006), I will assume suffix vowels are unspecified for backness when harmony occurs. We can then explain the variation in (44d) as the result of lexical selection of the suffix vowel.

The variable patterns occur when a back vowel is followed by one or more neutral vowels, shown in (45).

(45) Variable harmony with BN(N) stems

a. BN-F	[ɔrzeːn-nɛk]	*arzénnek*	'arsenic-DAT'
b. BN-B	[ɔrzeːn-nɔk]	*arzénnak*	'arsenic-DAT'

Hayes & Londe (2006)

This last case is the interesting one, as it is the only pattern that is not fully deterministic: for the same stem, the suffix can either harmonize or not harmonize, and both are acceptable. For the patterns in (44), I adopt Hayes and Londe's (2006) analysis, as I use the same constraints and there is no variation involved; here I will only provide an analysis of the variable pattern in (45). Table 4.1 summarizes the possible patterns.

Variation in the B(N)N cases is conditioned by two phenomena our analysis must account for. The first is the "height effect": the lower the (final) neutral vowel in the stem, the more likely the suffix vowel is to be [-back], with the relationship between neutral vowels shown in (46).

(46) The height effect

P([-back] | [ɛ]) > P([-back] | [eː]) > P([-back] | [i iː])

In fact, as Hayes and Londe (2006) determined (with a wug test), [i iː] hardly ever elicits a [-back] suffix vowel at all. This is not the case for [ɛ] and [eː], which both appear with either [+back] or [-back] suffix vowels at varying rates.

The other phenomenon is the "count effect": BNN stems are more likely to elicit [-back] than BN stems, no matter the height of the neutral vowel. Of course, these two phenomena are cumulative: for example, BNɛ stems are the most likely to be associated with a [-back] suffix vowel.

(47) The count effect

P([-back] | BN_1N_2) > P([-back] | BN_2)

Table 4.2 gives the frequencies reported in Hayes and Londe (2006) of [+back] suffix vowels for each stem type, demonstrating both the height and count effects. For any given final vowel, the frequency of [+back] is greater for the B(x) configuration than

TABLE 4.1. Summary of Hungarian vowel harmony

Stem vowel(s)	Suffix vowel
1. (x)B	[+back]
2. (x)F	[-back]
3. (N)N	Usually [-back] [+back] with select stems, e.g. *híd*
4. F(N)N	[-back]
5. B(N)N	Stem-specific choices (like *híd*) Free variation, e.g. *arzénnek/arzénnak*

for the BN(x) configuration, while the height effect is followed within each group of B(x) and BN(x) stems.

An analysis which uses general suppressible constraints, not relativized to a particular vowel height, to motivate harmony will not be able to account for the height effect (and, consequently, cannot account for the count effect relative to each neutral vowel), as all neutral vowels would be predicted to draw [+back] suffix vowels at the same rate. Consider the approach taken in Ringen and Vago (1998) and Kimper (2011b) (through Harmonic Serial Variation), where harmony is motivated by a constraint encouraging the spread of a [+back] feature from the stem to the suffix as in (48a).

(48a) Θ SPREAD[+BACK] ≡
Assign one violation for each [+back] feature that is linked to exactly one segment.

(48a), a version of the constraint adapted for MS, favors either the deletion of individual [+back] features or their absorption into a feature chain of some kind (Walker 2011; see Section 3.3.1 for more discussion of feature chains).

TABLE 4.2. Observed [+back] frequencies

Stem type	[+back] Frequency
1. Bi	0.953
2. Beː	0.376
3. Bɛ	0.071
4. BNi	0.287
5. BNeː	0.073
6. BNɛ	0.015

Competing with (48a) is another SPREAD constraint, this time encouraging the spread of [-back].

(48b) (Θ) SPREAD[-BACK] ≡
Assign one violation for each [-back] feature that is linked to exactly one segment.

Every output form will violate one of the constraints in (48). The competition between these two constraints will decide whether or not harmony occurs. What it will NOT do, however, is differentiate between neutral vowels and their different tendencies toward transparency, as required by the height effect. To achieve that, other constraints are needed. Since SPREAD[-BACK] blocks [+back] harmony, an obvious way to do this is to relativize this constraint to particular neutral vowels. Crucially, these constraints must not compete with (48a) on the level of the constraint hierarchy since the rankings between them may be unreliable.

This roughly describes the approach taken by Hayes and Londe (2006), and in the next section I will show that their analysis is portable to Markedness Suppression.

4.2.2 Analysis

As noted above, my analysis is based on Hayes and Londe (2006). Though their analysis used Stochastic OT in order to produce the correct frequencies, I will show that their constraints are also compatible with a Markedness Suppression approach to variation. Those constraints are given in (49).

(49) a. LOCAL[X] (*X* ∈ {*F*, *B*, *NN*, *i*, *ɛ*, *iː*, *eː*}) ≡
Assign one violation for any pair of adjacent vowels, one of which must be X, that does not share specifications for [back].

b. DISTAL[X] (*X* ∈ {*F*, *B*}) ≡
Assign one violation for any pair of vowels, one of which must be X, that does not share specifications for [back].

These are AGREE constraints, specific to [back] and a supplied argument (some class of vowel), that vary in sensitivity to locality. LOCAL[X] penalizes configurations where a vowel of class X is IMMEDIATELY followed or preceded by a vowel with a different backness specification (i.e. a vowel of a different class). DISTAL[X] differs only in that it is indifferent to the distance between the two vowels: for every vowel of class X, it will assign one violation for every other vowel in the word with a different backness specification.

In Hayes and Londe's (2006) analysis, the arguments of these constraints were either schematized classes of vowels (F, B, N, NN) or individual neutral vowels. However, since the probability of suppression for each of the constraints used in our analysis must be the same, this would lead to an analysis that is unable to account for the height effect, as all of the neutral vowels would vary in choice of suffix vowel at the same rate. This is shown in (50).

(50) Tableaux for Hungarian vowel harmony (no height effect)

I. /Bε-A/	Θ LOCAL[ε]	Θ LOCAL[eː]	Θ LOCAL[i]	DISTAL[B]
→ a. Bε-F				**
(→) b. Bε-B	*			*
II. /Beː-A/	Θ LOCAL[ε]	Θ LOCAL[eː]	Θ LOCAL[i]	DISTAL[B]
→ a. Beː-F				**
(→) b. Beː-B		*		*
III. /Bi-A/	Θ LOCAL[ε]	Θ LOCAL[eː]	Θ LOCAL[i]	DISTAL[B]
→ a. Bi-F				**
(→) b. Bi-B			*	*

In (50I), the frequency of the variant in (b) is equal to the probability of suppression of Θ LOCAL[ε]. Similarly, in (50II), the frequency of (b) is equal to the probability of suppression of Θ LOCAL[eː] and in (50III), the frequency of (b) is equal to the probability of suppression of Θ LOCAL[i]. However, all of these constraints have the SAME probability of suppression, as *p* is set on a per-language basis. Consequently, there is no height effect in this analysis.

Instead of referring to individual vowels, I will take advantage of the fact that requiring multiple violations to be suppressed reduces the probability of a given output form and use stringency (de Lacy 2003). Since the height effect forms a natural hierarchy, we can encode that hierarchy into our AGREE constraints by having constraints pertaining to vowels LOWER on the hierarchy assign violations to those HIGHER than them on the hierarchy as well. Since there are only three vowel heights in the hierarchy, we only need to adjust one constraint to achieve this.

(51) Θ LOCAL[ε eː] ≡
Assign one violation for any pair of adjacent vowels, one of which is either [ε] or [eː], that do not share specifications for [back].

This constraint requires local harmony for both [eː] and [ε]. To complete the pattern, we also need Θ LOCAL[N], which requires local harmony for all of the neutral vowels, and Θ LOCAL[ɛ], which only requires local harmony for [ɛ]. Since the height effect does not GUARANTEE victory for one variant over another in any of these cases, all three of these must be suppressible. The stringency relationship between these constraints is shown in Table 4.3.

In order for any of the stem-suffix configurations in Table 4.3 to surface, all of the violations assigned to that configuration by the LOCAL family of constraints would need to be suppressed; otherwise, the corresponding BN-F form will surface. Consequently, a stem-suffix combination like a that receives three violations total – one for each of Θ LOCAL[ɛ], Θ LOCAL[ɛ eː], and Θ LOCAL[N] – will only surface if all of the violations assigned by those three constraints are suppressed. On the other hand, a variant like (c) in Table 4.3 would only need to suppress one violation, the one given to it by Θ LOCAL[N], in order to surface. Hence, the form (c) will occur with frequency *p*, corresponding to a need for one violation suppressed, while the form a has the smaller frequency p^3, corresponding to a need for three violations to be suppressed.

These constraints DO interact, in that for two of them the environments they target form a subset of the environments targeted by at least one other constraint. Yet there is no crucial ranking between them, as they would all assign violations for the same marked

TABLE 4.3. Stringency relations of LOCAL constraints

Stem-suffix Configuration	Violates Θ LOCAL[ɛ]?	Violates Θ LOCAL[ɛ eː]?	Violates Θ LOCAL[N]?
a. Bɛ-B	Yes	Yes	Yes
b. Beː-B	No	Yes	Yes
c. Bi-B	No	No	Yes

feature at each locus. However, these must all crucially outrank DISTAL[B], as otherwise the suffixes would always be [+back].

The results for each possible BN stem are demonstrated in the tableaux in (52).

(52) Tableaux for the height effect

I. /Bɛ-A/	ʘ LOCAL[ɛ]	ʘ LOCAL[ɛ eː]	ʘ LOCAL[N]	DISTAL[B]
→ a. Bɛ-F				**
(→) b. Bɛ-B	*	*	*	*
II. /Beː-A/	ʘ LOCAL[ɛ]	ʘ LOCAL[ɛ eː]	ʘ LOCAL[N]	DISTAL[B]
→ a. Beː-F				**
(→) b. Beː-B		*	*	*
III. /Bi-A/	ʘ LOCAL[ɛ]	ʘ LOCAL[ɛ eː]	ʘ LOCAL[N]	DISTAL[B]
→ a. Bi-F				**
(→) b. Bi-B			*	*

In (52I), candidate (b) only wins if each of the violations it earns from each of the LOCAL constraints is suppressed. In (52II), candidate (b) does not get a violation from ʘ LOCAL[ɛ] (because it doesn't have an [ɛ]), so it only needs both the violations from ʘ LOCAL[ɛ eː] and ʘ LOCAL[N] to be suppressed in order to win. Finally, in (52III), candidate (b) wins whenever the most general constraint ʘ LOCAL[N] is suppressed, as it does not have any of the vowels subject to the more specific LOCAL constraints.

In order to account for the count effect, we can include a fourth LOCAL constraint: ʘ LOCAL[NN].[10] Since the count effect, like the height effect, is not absolute, this constraint is also suppressible. Again, there is no crucial ranking between this constraint and the other suppressible LOCAL constraints; however, it must outrank DISTAL[B] for the same reason as those did. This ranking is demonstrated in (53), which uses a BNeː wug

[10] This can also be accounted for with ʘ DISTAL[N], though that would also influence the number of violations that would need to be suppressed for the height effect, and consequently the modeled frequencies for those forms (essentially multiplying the rate of all the transparent forms by *p* due to the additional violation that needs to be suppressed in each case). Since Hayes and Londe (2006) use this constraint, I have opted to use it here as well.

stem from Hayes and Londe (2006).

(53) Tableau for /aːɲiveːl-nAk/ (wug stem) and the count effect

/aːɲiveːl-nAk/	Θ LOCAL[NN]	Θ LOCAL[ɛ eː]	Θ LOCAL[N]	DISTAL[B]
→ a. **aːɲiveː**l-nɛk BNNF				***
(→) b. **aːɲiveː**l-nɔk BNNB	*	*	*	**

In order for (b) to be optimal in (53), violations must be suppressed from all of the LOCAL constraints. Otherwise, (a) will be optimal, as any of the unsuppressed violations will kill (b) before DISTAL[B] can pick it over (a). This translates to a probability of p^3 for (b) and $1 - p^3$ for (a). This will be the case for all BNeː stems. For the others, the probabilities differ depending on how many LOCAL constraints will assign violations.

Tables 4.4 and 4.5 summarize the expected frequencies for all of these forms in terms of *p*. Table 4.4 gives the frequencies for BN stems only, while Table 4.5 gives frequencies for BNN stems, which are just the frequencies of the corresponding BN stems multiplied by *p*, corresponding to the extra violation of Θ LOCAL[NN] that would need to be suppressed for the [+back] forms to surface with these stems.

Note that, since $p < 1$, the greater the exponent of *p*, the less frequent the [+back] variant becomes. So, for example, Bɛ-B forms, with a frequency of p^3, as LESS frequent than Bi-B forms, with a frequency of *p*.

TABLE 4.4. BN frequencies

Final stem vowel	[+back] suffix frequency	[-back] suffix frequency
1. ɛ	p^3	$1 - p^3$
2. eː	p^2	$1 - p^2$
3. i	p	$1 - p$

TABLE 4.5. BNN frequencies

Final stem vowel	[+back] suffix frequency	[-back] suffix frequency
1. ɛ	p^4	$1 - p^4$
2. eː	p^3	$1 - p^3$
3. i	p^2	$1 - p^2$

4.2.3 Discussion

The analysis above is capable of modeling the existence of the height and count effects in general proportions: for example, BN_1N_2 stems lead to harmony more often than BN_2 stems. What it cannot do with any reasonable degree of accuracy is model the actual frequencies as given in and used by Hayes and Londe (2006). Using the Euclidean metric,[11] the value of *p* (which corresponds to frequencies of [+back] suffixes throughout the analysis) that most closely approximates all of the desired frequencies is approximately 0.585. This leads to the following frequencies for each category of stem given in Table 4.6 and represented graphically in Figures 4.1 and 4.2.

Only the model frequencies for Beː and BNi are less than 0.1 away from the actual frequencies; the model predicts harmony through [i] a troubling 36.8% more often than it actually happens. More importantly, the model frequencies paint a much different portrait of the overall pattern than the actual frequencies do. For those, there is an

[11] That is, I solved the problem

$$p^* = argmin\sqrt{\begin{array}{c}(p-0.953)^2+(p^2-0.376)^2\\+(p^3-0.071)^2+(p^2-0.287)^2+(p^3-0.073)^2+(p^4-0.015)^2\end{array}}.$$

Here, the minuends are the predicted rates for each output pattern in terms of *p* (see Tables 4.4 and 4.5), while the subtrahends are the attested output frequencies (Hayes & Londe 2006). Alternative means of calculating the closeness of the modeled frequencies to the actual frequencies will yield different values of *p*, with different degrees of accuracy in each of the categories modeled, but none of them will yield values of *p* that closely approximate all of the actual frequencies. This is clear from the problem above: for example, we need a *p* such that p^2 is close to *both* 0.376 *and* 0.287, numbers that are themselves not particularly close.

TABLE 4.6. Model vs. actual [+back] suffix frequencies

Stem type	Model [+back] frequency	Actual [+back] frequency	Error
1. Bi	0.585	0.953	0.368
2. Beː	0.342	0.376	0.034
3. Bɛ	0.200	0.071	0.129
4. BNi	0.342	0.287	0.055
5. BNeː	0.200	0.073	0.127
6. BNɛ	0.117	0.015	0.102

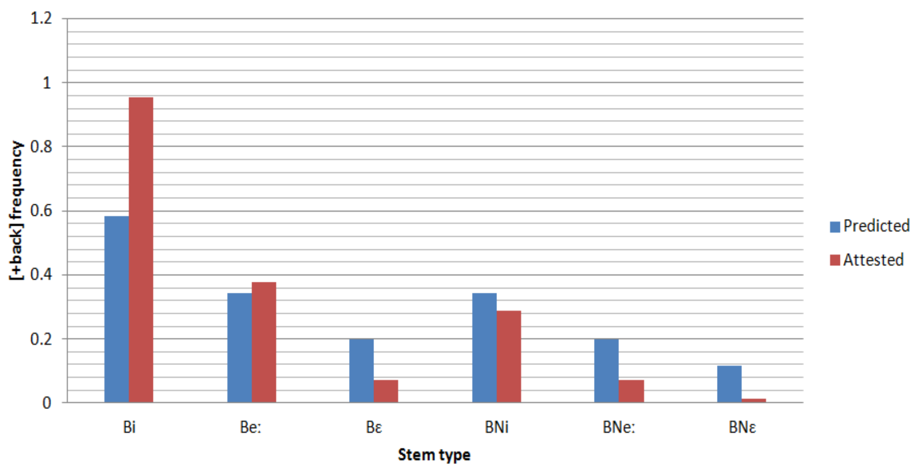

FIGURE 4.1. Model vs. actual [+back] suffix frequencies (comparison)

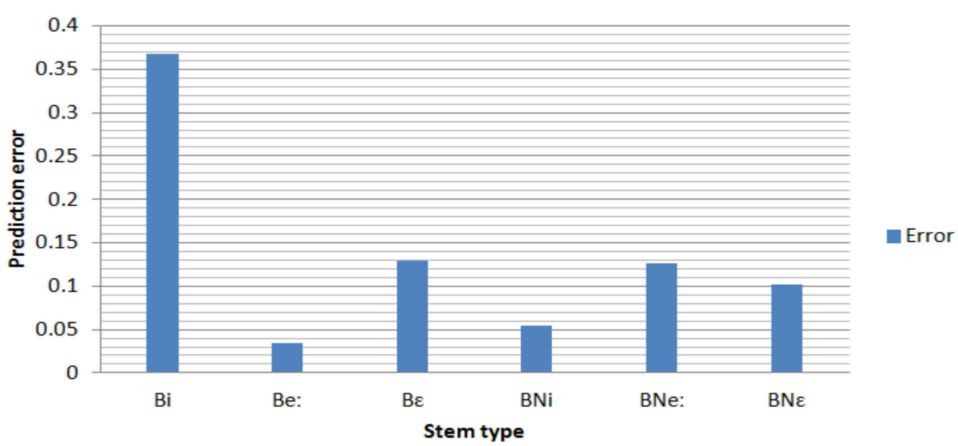

FIGURE 4.2. Model vs. actual [+back] suffix frequencies (absolute error)

enormous drop off between the stems with [i] and the other neutral vowels in both BN and BNN cases. In the BN case, it seems as though each neutral vowel patterns independently, with [i] almost always transparent and [ɛ] almost always not, while in the BNN case, the vowels [ɛ eː] seem to pattern together, independent of [i]. None of this is reflected in the model frequencies, which predict a strictly proportionate drop off of [+back] suffix vowels between each vowel along the dimension of height.

One solution to this problem would be to stratify LOCAL[NN] constraints, that is, introduce constraints like LOCAL[Nɛ Neː] to produce a second stringency relationship among the vowels of the BNN stem. This would impose the height effect into BNN stems directly. However, the count effect would be reduced to the result of a stem receiving violations from both the relevant BNN and BN LOCAL constraints, as opposed to this analysis, where the count effect is a result of a specific constraint, LOCAL[NN]. This strikes me as a case of overparametrization.

As a parting thought, it is worth pointing out that this section assumed, on account of the data, that each of the LOCAL constraints is suppressible, though that is not necessarily the case. Because these constraints all favor [-back] suffix vowels, if any of them is not suppressible, then all of the neutral vowels assigned violations by that constraint will uniformly select [-back] suffix vowels.

For example, if LOCAL[ɛ eː] were not suppressible, the outcome would resemble the tableaux in (54). This is an extreme version of the height effect. Now, only [i] can allow transparent harmony to a [+back] suffix vowel: the other two classes are always opaque. Though this does not correspond to the pattern in Hungarian, if the height effect, is in fact a natural hierarchy, it is reasonable to predict a language could implement it more stringently, as it is here, making Hungarian's statistical tendencies into categorical patterns.

(54) Tableaux for the height effect; LOCAL[ɛ eː] not suppressible

I. /Bɛ-A/	Θ LOCAL[ɛ]	LOCAL[ɛ eː]	Θ LOCAL[N]	DISTAL[B]
→ a. Bɛ-F				**
b. Bɛ-B	*	*!	*	*
II. /Beː-A/	Θ LOCAL[ɛ]	LOCAL[ɛ eː]	Θ LOCAL[N]	DISTAL[B]
→ a. Beː-F				**
b. Beː-B		*!	*	*
III. /Bi-A/	Θ LOCAL[ɛ]	LOCAL[ɛ eː]	Θ LOCAL[N]	DISTAL[B]
→ a. Bi-F				**
(→) b. Bi-B			*	*

On the other hand, if LOCAL[N] were not suppressible, the result would be opacity everywhere. This is shown in (55).

(55) Tableaux for the height effect; LOCAL[N] not suppressible

I. /Bɛ-A/	Θ LOCAL[ɛ]	Θ LOCAL[ɛ eː]	LOCAL[N]	DISTAL[B]
→ a. Bɛ-F				**
b. Bɛ-B	*	*	*!	*
II. /Beː-A/	Θ LOCAL[ɛ]	Θ LOCAL[ɛ eː]	LOCAL[N]	DISTAL[B]
→ a. Beː-F				**
b. Beː-B		*	*!	*
III. /Bi-A/	Θ LOCAL[ɛ]	Θ LOCAL[ɛ eː]	LOCAL[N]	DISTAL[B]
→ a. Bi-F				**
b. Bi-B			*!	*

In effect, this produces a system where the neutral vowels are no longer neutral, consistently harmonizing to [-back] as if they were a part of the F class of vowels. This is, again, not the pattern in Hungarian, but as it is simply a different kind of vowel harmony, this prediction is also acceptable. Similarly, this analysis has assumed that DISTAL is not suppressible. If it were, it would increase the number of scenarios in which opacity wins across the board, diminishing the neutrality of these vowels, without inverting the height effect.[12]

To conclude, at the beginning of this section, I gave an example of an analysis of Hungarian vowel harmony that relied on constraint domination and therefore could not work under Markedness Suppression. In this analysis, the only variable constraint was the one that enforced harmony generally; those that affected its application in certain environments were not suppressible. I then showed that it was possible to produce a working model of this pattern using Markedness Suppression if instead of using a single suppressible constraint, the requirement of harmony were handled by three suppressible separate constraints, each specific to one of the vowels of the height effect, as well as another constraint for the count effect. In order to make frequency predictions according to the height and count effects, stringency was used to make sure some output patterns occurred less often. Though this was successful in implementing the general predictions of the height and count effects, it was not successful when it came to modeling corpus frequencies.

4.3 English t/d deletion

4.3.1 Introduction

In many dialects of English – Philadelphians and New Yorkers (Guy 1980) and Chicanos (Santa Ana 1991), among others – the phones [t d] are subject to deletion in the environment C_#. Deletion is variable, conditioned on what follows: a consonant, a

[12] In particular, the frequency of Bi-B must be greater than or equal to the frequency of Beː-B. If any of the LOCAL constraints is not suppressible, the frequency of the corresponding -B forms is 0, as shown above. This leaves only the case where all of the LOCAL constraints are suppressible. Including ties, Bi-B then surfaces if Ο LOCAL[N] is suppressed *and* Ο DISTAL[B] is suppressed (retained) in its favor, while Beː-B only surfaces if both Ο LOCAL[N] and Ο LOCAL[ɛ eː] are suppressed *and* Ο DISTAL[B] is suppressed (retained) in its favor. Since Ο DISTAL[B] assigns the same pattern of violations in both cases – that is, two violations to the -F form and one to the -B form – this means the frequency of Beː-B will be less than the frequency of Bi-B by a factor of *p*, owing to the requirement that an additional violation from Ο LOCAL[ɛ eː] be suppressed before Ο DISTAL[B] can be relevant.

vowel, or a pause. For example, in the Philadelphian dialect, we see the patterns in Table 4.7.

Across all dialects, the preconsonantal environment is most likely to elicit deletion (Kiparsky 1993), with certain consonants more likely than others to cause deletion: for example, [l] triggers deletion more often than [r] (Guy 1991, cited in Kiparsky 1993). Dialects differ with respect to frequency of deletion before pauses and vowels. Pauses, in particular, usually group with either consonants or vowels in frequency, although they can also admit the lowest frequency of deletion, depending on dialect (Kiparsky 1993). For example, in African American English in Washington, D.C., we observe the frequencies in Table 4.8 (Labov, Cohen, Robbins, & Lewis 1968).

In African American English in Table 4.8, deletion patterns with pre-consonant deletion. The opposite pattern is observed in Philadelphian English, as well as in Chicano English (Santa Ana 1991), as shown in Table 4.9. Finally, in New York English, the prepausal context is least likely to result in deletion, as shown in Table 4.10. Each dialect

TABLE 4.7. t/d Deletion in Philadelphian English

Context	Frequency of [t d] deletion
wɛs[t]# C ~ wɛs# C	100%
wɛs[t]# V ~ wɛs# V	12%
wɛs[t]# # ~ wɛs# #	38%

TABLE 4.8. t/d Deletion in African American English (Washington, D.C.)

Context	Frequency of [t d] deletion
wɛs[t]# C ~ wɛs# C	76%
wɛs[t]# V ~ wɛs# V	29%
wɛs[t]# # ~ wɛs# #	73%

TABLE 4.9. t/d Deletion in Chicano English

Context	Frequency of [t d] deletion
wɛs[t]# C ~ wɛs# C	61%
wɛs[t]# V ~ wɛs# V	32%
wɛs[t]# # ~ wɛs# #	33%

TABLE 4.10. t/d Deletion in New York English

Context	Frequency of [t d] deletion
wɛs[t]# C ~ wɛs# C	100%
wɛs[t]# V ~ wɛs# V	83%
wɛs[t]# # ~ wɛs# #	66%

brings with it a different frequency of deletion in each environment and, thus, a different set of proportions that must be mirrored by the frequencies projected by the analysis. Additionally, Guy (1991) and Santa Ana (1991) observed an interaction with the strength of morpheme boundaries and deletion: the past tense suffix [t] or [d] often creates the environment for deletion, but the probability of deletion in this morphologically rich environment is less so than if deletion were all stem internal. In particular, if the probability of [t] deletion in a word like *cost* is p, then the probability of deletion in a word like *lost* (*lose+t*) is approximately p^2 and the probability of deletion in a word like *toss+ed* is approximately p^3 (Kiparsky 1993). The generalization is that the less information a word-final [t d] carries, the more expendable it is, with the intermediate case being a stem-changing fusion (as in *lost*). This pattern is summarized in Table 4.11.

Following this chart, if the frequency of [t d] deletion occurring in the _#C environment is p = 61%, as in Chicano English, then we would expect the frequency of the [t] being deleted in *cost* to be approximately 61%, approximately p^2 = 37% in *lost*,

TABLE 4.11. Influence of morphophonological context on [t d] deletion

Word-internal Context	Frequency of [t d] deletion (*p* = frequency of deletion in some environment)
cost	p
lost (lose + t)	p^2
toss + ed	p^3

and approximately p^3 = 23% in *tossed*. These estimates accurately summarize the data (Kiparsky 1993, Santa Ana 1991).

Setting aside the question of frequency modeling, like Hungarian vowel harmony, whether MS can handle this system at all depends on the analysis chosen. Kiparsky (1993), for example, uses a set of constraints that, although amenable to a partial orders analysis, do not lend themselves to a Markedness Suppression reanalysis. His approach to this phenomenon is based on the ability or inability to resyllabify across word boundaries in order to avoid violating word-final phonotactics, exemplified by the constraint in (56).

(56) SYLLABLE-WELLFORMEDNESS (SYLLWF) ≡
Equivalent to NOCODA and *COMPLEX.

This is the markedness constraint used to encourage deletion in the first place by penalizing complex codas.[13] It also prevents resyllabification to the following word if that word begins with a consonant, as that would create another consonant cluster; resyllabification can only be used as a repair strategy to avoid violating this constraint if the following syllable has no onset.

(57) ALIGN ≡
No resyllabification across word boundaries and no deletion of phrase-final consonants.

[13] Strictly speaking, according to Kiparsky (1993), this constraint would penalize all codas; however, since t/d-deletion only occurs in the context of C[t d]# codas, it is more reasonable to interpret it as a constraint penalizing only complex codas. Otherwise, something would have to happen to the pre-[t d] C as well.

This constraint penalizes attempts to repair illicit C[t d]# by resyllabifying, while preserving word-final [t] or [d] in prepausal position. Its purpose is to favor preserving [t] or [d] in prepausal position while remaining indifferent to deletion everywhere else.

(58) PARSE ≡
No unrealized segments.

Finally, (58) is Kiparsky's (1993) approach to what is now more generally MAX. Accordingly, its role is to prevent deletion everywhere.

In Kiparsky's (1993) analysis, the dominance relationships between these three constraints determined which syllabification options are available and, as a result, whether deletion would occur in a given environment. The possible patterns, and the rankings that produce them in a Partial Orders analysis, are given in tableaux (59-62).[14]

In (59), since PARSE eliminates all of the deletion candidates, word-final [t] or [d] will be preserved in some way under this ranking. The exact mechanism by which this is achieved – resyllabification or not – depends on the actual ranking between SYLLWF and ALIGN, which I left undecided in (59).

Both SYLLWF and ALIGN penalize at least one of the preservation methods in tableaux (60I,II), killing them, and SYLLWF in particular kills the faithful candidate in tableau (60III) before ALIGN can save it. Consequently, deletion happens everywhere.

The main difference between the tableaux in (61) and those in (60) are that now ALIGN eliminates the deletion candidate in tableau (61III) before SYLLWF kills it so that retention is preferred prepausally. Deletion still occurs in the other two environments, for the same reason it did in (60).

[14] A fifth pattern, deletion before consonants only, is possible under Kiparsky's (1993) analysis, but requires the ALIGN constraint to be split into its component parts. As the present analysis's failure on the remainder of the data is sufficient to reject it, this is unnecessary for the present discussion and has been left out.

In (62), high-ranking SYLLWF eliminates the faithful candidates in the preconsonantal and prepausal contexts, but because SYLLWF does not assign violations when resyllabification occurs in the prevocalic environment and PARSE outranks ALIGN, the constraint that DOES, the result is deletion everywhere except before vowels.

If we used exactly these constraints to produce a Markedness Suppression analysis of t/d-deletion, we could replicate the empirical observation that deletion is variable in broad strokes, but we would be unable to precisely replicate the observed relationships between those frequencies for most dialects. Since PARSE, a faithfulness constraint, cannot be suppressible and because of the broadness with which SYLLWF and ALIGN are defined, each calling for deletion in multiple environments, there are only two possible rankings that could yield any results on par with the observed trends: one for dialects in which deletion in the _#C environment is categorical and another for dialects in which that is not the case. This turns out to be far too limiting, as I will now show.

(59) Deletion nowhere: PARSE >> SYLLWF

I. /wɛst C/	PARSE	SYLLWF	ALIGN
→ a. wɛst] [C		*	
b. wɛs] [tC		*	*!
c. wɛs] ŧ [C	*!		
II. /wɛst V/	PARSE	SYLLWF	ALIGN
→ a. wɛst] [V		*	
→ b. wɛs] [tV			*
c. wɛs] ŧ [V	*!		
III. /wɛst #/	PARSE	SYLLWF	ALIGN
→ a. wɛst] #		*	
b. wɛs] ŧ #	*!		*

(60) Deletion everywhere: SYLLWF >> ALIGN >> PARSE

I. /wɛst C/	SYLLWF	ALIGN	PARSE
a. wɛst] [C	*!		
b. wɛs] [tC	*!	*	
→ c. wɛs] ɬ [C			*
II. /wɛst V/	SYLLWF	ALIGN	PARSE
a. wɛst] [V	*!		
b. wɛs] [tV		*!	
→ c. wɛs] ɬ [V			*
III. /wɛst #/	SYLLWF	ALIGN	PARSE
a. wɛst] #	*!		
→ b. wɛs] ɬ #		*	*

(61) Deletion before C and V: ALIGN >> SYLLWF >> PARSE

I. /wɛst C/	ALIGN	SYLLWF	PARSE
a. wɛst] [C		*!	
b. wɛs] [tC	*!	*	
→ c. wɛs] ɬ [C			*
II. /wɛst V/	ALIGN	SYLLWF	PARSE
a. wɛst] [V		*!	
b. wɛs] [tV	*!		
→ c. wɛs] ɬ [V			*
III. /wɛst #/	ALIGN	SYLLWF	PARSE
→ a. wɛst] #		*	
b. wɛs] ɬ #	*!		*

(62) Deletion before C and #: SYLLWF >> PARSE >> ALIGN

I. /wɛst C/	SYLLWF	PARSE	ALIGN
a. wɛst] [C	*!		
b. wɛs] [tC	*!		*
→ c. wɛs] ɬ [C		*	
II. /wɛst V/	SYLLWF	PARSE	ALIGN
a. wɛst] [V	*!		
→ b. wɛs] [tV			*
c. wɛs] ɬ [V		*!	
III. /wɛst #/	SYLLWF	PARSE	ALIGN
a. wɛst] #	*!		
→ b. wɛs] ɬ #		*	*

First, consider dialects in which deletion in the _#C environment as categorical. Above, I referred to New York and Philadelphian English as examples of this pattern; the frequencies of deletion for these dialects are in Tables 4.7 and 4.10.

In addition to showing a considerable difference in the willingness to delete in the other environments at all, these dialects are mirror images of each other with respect to the frequency of deletion in the non-_#C contexts: in New York English, deletion in the _#V environment is more likely than in the _## environment, and vice-versa in Philadelphian English. Ideally, we could capture these alternative preferences in our model for deletion in these dialects, but that is not possible with Kiparsky's (1993) constraints.

The constraint that drives deletion in the _#C environment is SYLLWF, so it cannot be suppressible if deletion is categorical, and we must have SYLLWF >> PARSE if the constraint is to have any effect at all. This suggests ALIGN must be suppressible in order to allow for noncategorical deletion in the other environments. Since deletion is not categorical elsewhere, ALIGN must outrank SYLLWF, or we would simply get the pattern shown in (60). Thus, in any dialect where deletion is categorical in the _#C environment, we must have the ranking in (63).

(63) Θ ALIGN >> SYLLWF >> PARSE

As shown in the tableau below, this successfully predicts categorical deletion in the _#C environment and gradient deletion elsewhere – but it also predicts this deletion will occur at the same rate in both the _#V and _## environments since in each case, only the suppression or nonsuppression of a single violation of Θ ALIGN decides between deletion and retention.

(64) Tableaux for categorical deletion in the _#C environment

I. /wɛst C/	Θ ALIGN	SYLLWF	PARSE
a. wɛst] [C		*!	
b. wɛs] [tC	*	*!	
→ c. wɛs] ŧ [C			*
II. /wɛst V/	Θ ALIGN	SYLLWF	PARSE
a. wɛst] [V		*!	
(→) b. wɛs] [tV	*		
→ c. wɛs] ŧ [V			*
III. /wɛst #/	Θ ALIGN	SYLLWF	PARSE
→ a. wɛst] #		*	
(→) b. wɛs] ŧ #	*		*

This both fails to adequately capture any dialect, like New York or Philadelphian English, in which there is a difference between the rates of deletion in the _#V and _## environments alongside categorical deletion in the _#C environment, as well as the differences between those dialects.

The situation is not improved when considering dialects in which deletion is never categorical. The examples of this pattern of deletion given above were Washington, D.C., African American English and Chicano English; the frequency distributions for each of these dialects are given in Tables 4.8 and 4.9.

As in New York and Philadelphian English, there are clear differences in how each dialect treats each environment: in African American English, deletion in the _## environment is as common as deletion in the _#C environment, with deletion least likely in the _#V environment, but in Chicano English, the _#V and _## environments pattern together. Any model of these frequencies should be able to show that, but that is, again, not possible with the constraints used here.

Both ALIGN and SYLLWF require categorical behavior – deletion in the _#C and _## environments for SYLLWF, retention in the _## environment for ALIGN – so both of them must be suppressible if there is no categorical deletion in the dialect. Since PARSE >> (Θ) ALIGN categorically mandates RETENTION in the _#V environment (cf. (60)) and as PARSE >> (Θ) SYLLWF would prevent deletion everywhere (cf. (59)), the only possible rankings for any dialect where deletion occurs in every environment but is never categorical are given in (65).

(65) a. Θ ALIGN >> Θ SYLLWF >> PARSE
b. Θ SYLLWF >> Θ ALIGN >> PARSE

Both of these result in a ranking that permits optional deletion everywhere, as shown in tableaux (66) and (67). The only difference between these two rankings arises in the violation profile in the _#C environment, which controls the model probabilities for that environment (as well as the model probabilities overall if all three environments are considered simultaneously for these projections, as I have done).

Because the pattern of violations in this case is roughly the same for each dialect, the only meaningful way to vary the rate at which deletion occurs in each environment in order to model differences between dialects is to change the rate of deletion, but the proportional relationships between the rates of deletion remain the same. This is clearest in the Θ SYLLWF >> Θ ALIGN ranking, where the total probability of deletion in the _#C environment is the same as the total probability of deletion in the _#V environment.[15]

[15] Referring to tableaux (67I), deletion occurs in the _#C environment only when no violation is suppressed (probability $(1 - p)^3$) or the only violation suppressed is ALIGN. For candidate (b) (probability $p(1 - p)^2$), generating a total probability of $(1 - p)^3 + p(1 - p)^2$, which simplifies to $(1 - p)^2$. Meanwhile, deletion occurs in the _#V environment only if neither violation in tableau (67II) is suppressed, which is associated with a probability of $(1 - p)^2$, or exactly the probability of deletion in the _#C environment.

(66) Tableaux for no categorical deletion (Θ ALIGN >> Θ SYLLWF >> PARSE)

I. /wɛst C/	Θ ALIGN	Θ SYLLWF	PARSE
(→) a. wɛst] [C		*	
(→) b. wɛs] [tC	*	*	
→ c. wɛs] ɬ [C			*
II. /wɛst V/	Θ ALIGN	Θ SYLLWF	PARSE
(→) a. wɛst] [V		*	
(→) b. wɛs] [tV	*		
→ c. wɛs] ɬ [V			*
III. /wɛst #/	Θ ALIGN	Θ SYLLWF	PARSE
→ a. wɛst] #		*	
(→) b. wɛs] ɬ #	*		*

(67) Tableaux for no categorical deletion (Θ SYLLWF >> Θ ALIGN >> PARSE)

I. /wɛst C/	Θ SYLLWF	Θ ALIGN	PARSE
(→) a. wɛst] [C	*		
(→) b. wɛs] [tC	*	*	
→ c. wɛs] ɬ [C			*
II. /wɛst V/	Θ SYLLWF	Θ ALIGN	PARSE
(→) a. wɛst] [V	*		
(→) b. wɛs] [tV		*	
→ c. wɛs] ɬ [V			*
III. /wɛst #/	Θ SYLLWF	Θ ALIGN	PARSE
→ a. wɛst] #	*		
(→) b. wɛs] ɬ #		*	*

These rates will always be identical, no matter the choice of p. Using the technique for generating the closest approximation to empirical frequencies used for the Hungarian data[16] for both rankings, we get the model frequencies for the Θ SYLLWF >>

[16] In particular, I solved the problem

$$p^* = argmin \sqrt{\left(\left([1-p]^3 + 2p[1-p]^2\right) - \#C\right)^2 + \left((1-p)^2 - \#V\right)^2 + \left(p(1-p) - \#\#\right)^2}$$

for the ranking Θ ALIGN >> Θ SYLLWF, and

$$p^* = argmin \sqrt{\left((1-p)^2 - \#C\right)^2 + \left((1-p)^2 - \#V\right)^2 + \left(p(1-p) - \#\#\right)^2}$$

for the ranking Θ SYLLWF >> Θ ALIGN, where $\#C$ refers to the probability of deletion in the _#C environment for a given dialect, and similar for $\#V$ and $\#\#$, while the expressions with p are the total probability of deletion in each environment, derived by taking the sum of the probabilities for each possible configuration of retained/suppressed violations that would lead to deletion in that environment. For

ϴ ALIGN ranking shown in Table 4.12. The pattern in which the _#C and _#V environments pattern so closely they might as well be identical, while deletion in the _## environment is much lower, is both unattested and the only plausible outcome of this ranking.[17]

The closest example is New York English, where deletion in the _#C environment is categorical and deletion in both of the other environments occurs over 60% of the time. This ranking, then, can likely be dismissed.

A somewhat better result is seen in the ϴ ALIGN >> ϴ SYLLWF ranking shown in Table 4.13. Though this does not produce a pattern we would not expect to see in SOME dialect, it not only fails to replicate the pattern of either of the two dialects it was to model; it also fails to distinguish between them. Though alternative methods of calculating *p* can grant more fidelity to the rate of deletion in one environment at the cost of greater error in the other environments, such that it is possible to grant more accuracy to one environment in one dialect and a different environment in another dialect so as to

TABLE 4.12. Model frequencies for African American and Chicano English (ϴ SYLLWF >> ϴ ALIGN)

Context	Predicted rate of [t d] deletion	
	African American (p = 0.323)	Chicano (p = 0.328)
wɛs[t]# C ~ wɛs# C	46%	45%
wɛs[t]# V ~ wɛs# V	46%	45%
wɛs[t]# # ~ wɛs# #	22%	22%

example, deletion in the _#V environment only occurs if neither the violation of ϴ ALIGN nor the violation of ϴ SYLLWF is suppressed, so deletion has a probability of $(1 – p)^2$ in this environment. Crucially, while the empirical frequencies depend on the dialect, the probabilities of deletion do not, as there is only one ranking that allows for optional deletion everywhere.

[17] It is also possible for deletion to be *more* likely in the _## environment, but this is contrary to the overall generalization that deletion is never more likely than in the _#C environment.

TABLE 4.13. Model frequencies for African American and Chicano English (Θ ALIGN >> Θ SYLLWF)

Context	Predicted rate of [t d] deletion	
	African American (*p* = 0.364)	Chicano (*p* = 0.379)
wɛs[t]# C ~ wɛs# C	55%	53%
wɛs[t]# V ~ wɛs# V	40%	39%
wɛs[t]# # ~ wɛs# #	23%	24%

produce a difference in the predictions between the two dialects, this would simply mask the underlying erroneous prediction that the proportional relationships between the rates of deletion in each environment should be the same, or nearly the same, for all dialects.

In summary, it is possible to produce the relevant output patterns in a Markedness Suppression analysis using Kiparsky (1993)'s constraints, but it is not possible to accurately show how each dialect differs in the RATE of deletion in each environment, making this analysis suboptimal.

The cause of this problem is the constraints themselves, which do too much and so limit the possible rankings, as well as the architecture of Markedness Suppression, which constrains the possible rankings by partially negating the dominance relationship between suppressible constraints and linking model probabilities more broadly to violation profiles.

In the following sections, I will show that it is possible to achieve far better empirical coverage by using an alternative analysis where each constraint references a particular environment, a strategy that avoids both of the obstacles imposed by Kiparsky's (1993) constraints.

4.3.2 Analysis

We must account for both the influence of the phonetic environment on deletion and the influence of the morphophonological circumstances. First, I will tackle the phonetic influences. Coetzee (2004) observes that the different environments in which deletion occurs correspond to a difference in cues for perceiving word-final [t d]: i.e. differences in the availability of a release burst or formant transitions conducive to the salience of a [t] or [d]. Table 4.14 gives his breakdown for how each environment fares in providing either of those cues.

Following Steriade (1997), Coetzee (2004) then proposes the positional markedness constraints in (68), motivated by the aforementioned differences in cues.

(68) *C[t d]#C, *C[t d]#V, *C[t d]##

Steriade (1997) suggests that contrastive features, which in this case includes word-final stops, are licensed by the cues in their environment – especially stops, which have no internal cues – and consequently are more or less likely to be preserved depending on the strengths of those cues. The constraints in (68) formalize this by providing a separate markedness constraint for each environment. Coetzee (2004) further suggests the universal ranking in (69) on the basis of the cues summarized in Table 4.14.

TABLE 4.14. Cues for [t d]# by context

Context	Release Burst	Formant Transitions
_#C	No[18]	Weakly before sonorants across a word boundary
_#V	Yes, across a word boundary	Yes, across a word boundary
_##	Yes	No

[18] As Coetzee (2004) points out, stops are regularly unreleased in this environment in English.

(69) *C[t d]#C >> *C[t d]#V, *C[t d]##

This amounts to a claim that *C[t d]#C is the most marked environment, reflecting the cross-dialectal observation that the _#C environment elicits deletion at least as often as either of the other two. With respect to the relationship between the _#V and _## environments, Coetzee points out that although the _#V environment is the best overall bearer of cues, those cues must be realized across a word boundary, whereas even though the _## environment does not provide formant transitions, the cue it DOES provide – the release burst – does not need to be realized across a word boundary. Consequently, there is "more freedom in how likely these two contexts are to sponsor a [t, d]" (Coetzee 2004: 228).

Unfortunately, even with the ranking in (69), a Markedness Suppression analysis on the basis of these constraints cannot reflect that generalization so long as *C[t d]#C is suppressible, as shown in (70).

(70) Tableaux for *west* (Coetzee's (2004) constraints)

I. /wɛst C/	ʘ *C[t d]#C	(ʘ) *C[t d]#V	(ʘ) *C[t d]##	MAX
(→) a. wɛst C	*			
→ b. wɛs C				*
II. /wɛst V/	ʘ *C[t d]#C	(ʘ) *C[t d]#V	(ʘ) *C[t d]##	MAX
(→) a. wɛst V		*		
→ b. wɛs V				*

In tableau (70I), candidate (a) survives so long as the lone violation for ʘ *C[t d]#C is suppressed, so its model frequency is equal to *p*. In tableau (70II), candidate (a)'s chances for survival depend on whether *C[t d]#V is suppressible: if it is not, then (a) can never win, while otherwise its chances are also *p*. A similar story holds for the _# environment. In order to maintain the generalization that _#C will never cause deletion less often than the other environments, we must assume that *C[t d]#V and *C[t d]## are always suppressible when *C[t d]#C is also suppressible, which means that this analysis always predicts deletion at the same rate for all environments, as there is only one violation assigned per constraint.

Luckily, there is an ordered relationship among each environment's propensity to cause deletion reflected in the natural ranking of constraints in (69). This provides a basis for a stringency relationship between these constraints, just as vowel height provided such a relationship for Hungarian, which will permit the analysis to guarantee that some output patterns are more likely than others. Here, it appears to be the strength of the cues for word-final [t d]: _#C is clearly the most marked environment cross-dialectally, followed by _#V in some dialects and _## in others.

(71) Hierarchy of cues
_#V, _## > _#C

As explained above, the relationship between the _## and _#V environments is not clear-cut due to the reliance of the latter on realizing cues across boundaries and the fewer cues available to the former. This shows up in the data, as well: in the African American English of Washington, D.C., the _#V environment is preferred, while the opposite is true in New York English and the two pattern together in Chicano English. I will treat the ambiguity in the preference for one or another of these environments here as a reflection of three possible hierarchies, each with its own stringency relationship.

(72) Unambiguous hierarchies
a. _#V > _## > _#C
b. _## > _#V > _#C
c. _## = _#V > _#C

Each dialect will then adhere to one of the hierarchies in (72), with its corresponding stringency relationship, and ignore the others. This analysis assumes, of course, that difference preferences in the data can be traced to different interpretation/valuation of cues; it is also possible that the _#V environment is unambiguously better at transmitting cues, and the empirical differences are due to something else. I will consider some alternatives to this description in the next section.

In the meantime, these hierarchies can be formalized with a constraint that penalizes environments with insufficiently strong cues.

(73) CueC[t d]#>X ≡
Assign one violation if the cues for [t d]# are worse than the cues for _X.

This produces three constraints, one for each environment.

(74) CueC[t d]#>C, CueC[t d]#>V, CueC[t d]#>#

Because the relationship between the _#C environment and the other two is the same in each dialect, the stringency relationship for CueC[t d]#>C is the same in every dialect. As discussed above, the relationship between V and # is not so consistent (although Kiparsky (1993) notes that # patterns closer to C more often, a fact not captured by this analysis). Consequently, those dialects in which _#V is more marked will have CueC[t d]#># assign violations to both _## and _#V environments and CueC[t d]#>V assign violations only to _## environments, while the reverse is true for languages where _## is more marked. In dialects where they are roughly on the same level, there is no stringency relationship between these two constraints.

The tableaux in (75) show one possible analysis using the constraints in (74). Here, all of those constraints are suppressible (though other optionality settings are possible; see Section 3.2 for discussion), as they would need to be for a dialect like Chicano or African American English in which deletion is never guaranteed in any environment. It is assumed that _#V is higher than _## on the hierarchy of cues for this dialect so that Θ CUEC[t d]#>V assigns violations to both environments. This corresponds to, among others, the African American dialect.

(75) Tableaux for *west* (final)

I. /wɛst C/	Θ CUEC[t d]#>C	Θ CUEC[t d]#>;#	Θ CUEC[t d]#>V	MAX
→ a. wɛst C	*	*	*	
→ b. wɛs C				*
II. /wɛst #/	Θ CUEC[t d]#>C	Θ CUEC[t d]#>#	Θ CUEC[t d]#>V	MAX
→ a. wɛst #		*	*	
→ b. wɛs #				*
III. /wɛst V/	Θ CUEC[t d]#>C	Θ CUEC[t d]#>#	Θ CUEC[t d]#>V	MAX
→ a. wɛst V			*	
→ b. wɛs V				*

The chief difference between (75) and (70) is that now _#C receives violations from all three of the suppressible constraints; all three violations must be suppressed in order to prevent deletion in (75I). In the other two environments, no more than two violation needs to be suppressed. This translates to a probability of retention of p^3 for _#C, p^2 for _##, and p for _#V, making _#C the most likely to delete by a factor of p. The exact value for p will, of course, depend on the data being modeled.

What if one of the CUEC[t d]# constraints is not suppressible? In that case, deletion always happens in every environment to which it assigns a violation. For the stringency relationship reflected in (75), that means that if CUEC[t d]#>V is not suppressible, deletion occurs everywhere; or if CUEC[t d]#># is not suppressible, deletion always occurs in the _#C and _## environments only; or if CUEC[t d]#>C is not suppressible, deletion always occurs only in the _#C environment. So long as CUEC[t d]#>C does not assign violations to any of the other two environments – and it cannot, given the hierarchy of cues in (72) – this system will never predict the unattested pattern of more deletion in the other two environments. Like Coetzee (2004), this analysis also predicts that it is not possible to have variable deletion in the _#C environment and categorical retention in the other environments since categorical retention would require $p = 1$, or the absence of all CUE violations, resulting in categorical retention in every environment due to MAX.

Recall, also, that morphology has an influence on deletion: if the word-final [t d] carries the information from an inflectional morpheme, it is less likely to delete. This follows a fairly straightforward pattern: if the probability of [t] deletion in a word like *cost* is p, then the probability of deletion in a word like *lost* (*lose+t*) is approximately p^2 and the probability of deletion in a word like *toss+ed* is approximately p^3 (Kiparsky 1993).

In order to account for the influence of morphology, we can split up each of the constraints in (74) into constraints sensitive to the type of boundary at the locus of violation, as Kiparsky (1993) does in his analysis.[19] For example, the constraint Θ CUEC[t d]#>V can be divided as follows, in order from smallest to largest domain.

(76) Θ CUEC[t d]#>V_{root} , Θ CUEC[t d]#>V_{stem} , Θ CUEC[t d]#>V_{word}

The ROOT-indexed constraint only assigns violations to root-internal clusters, corresponding to a word like *cost* or *west*. The STEM-indexed constraint assigns violations to clusters that are the result of a fusion of morphemes, like *lost*, as well as to root-internal clusters. Finally, the WORD-indexed constraint assigns violations regardless of the type of boundary. The result of this for the _#V environment can be seen in the tableaux

[19] Θ REALIZEMORPHEME, defined as a constraint that forces the realization of the morpheme [t d], would also achieve the desired result, but it would mess up the probabilities. We want a p, p^2, p^3 progression of frequencies for each of the morphophonological categories, but Θ REALIZEMORPHEME would only contribute p once. Even if it contributed p twice, that would lead to a progression of frequencies, for example, p^3, p^4, p^5 for the _#C environment, which represents the wrong proportional relationship.

in (77).

(77) Tableaux for morphophonology

I. /wɛst V/	Θ CUEC[t d]#>V$_{root}$	Θ CUEC[t d]#>V$_{stem}$	Θ CUEC[t d]#>V$_{word}$	MAX
→ a. wɛst V	*	*	*	
→ b. wɛs V				*
II. /lɔst V/	Θ CUEC[t d]#>V$_{root}$	Θ CUEC[t d]#>V$_{stem}$	Θ CUEC[t d]#>V$_{word}$	MAX
→ a. lɔst V		*	*	
→ b. lɔs V				*
III. /tɔs+t #/	Θ CUEC[t d]#>V$_{root}$	Θ CUEC[t d]#>V$_{stem}$	Θ CUEC[t d]#>V$_{word}$	MAX
→ a. tɔst #			*	
→ b. tɔs #				*

The probability of deletion of [t] in (77III) is simply $1 - p$, corresponding to the retention of a violation of Θ CUEC[t d]#>V$_{word}$. Similarly, for (77II), it is $(1 - p)^2$, and for (77I), it is $(1 - p)^3$. This follows precisely the exponential pattern of probabilities estimated by Guy (1991) and Santa Ana (1991). It is worth noting, though, that this analysis gives fairly extreme values. For one, the probability of deletion in the case of root-internal [t d] is $(1 - p)^9$ in the _#C environment since each of the violations from the nine Θ CUEC[t d]#$_Y$ constraints for all three phonetic environments must be retained in order for that to happen. If the probability of RETENTION in the least-marked _#V environment at the word level is a conservative 0.74, as it might be[20] in Chicano English (Santa Ana 1991), the probability of RETENTION at the root level in the most-marked _#C environment is then less than 0.07. Unfortunately, there are not enough data to determine whether this is TOO extreme a deviation.

[20] The data on deletion in each morphophonological environment are not partitioned by phonetic environment, so it is not clear how this value, which is the overall probability of retention at the word level, breaks down for _#C, et cetera.

4.3.3 Discussion

The model developed here is capable of showing the correct general order of markedness among the three phonetic environments – _#C, _#V, and _## – and is also capable of modeling the relationship among the three morphophonological environments – root, stem, and word – according to the approximately attested relationship between them. Despite its success in morphophonology, it is no more capable at handling the cross-dialectal frequency of deletion in each context than the analysis of Hungarian is at handling the height effect.

Table 4.15 gives the frequencies of retention by context for each of the four dialects. These four are broadly representative of the kinds of patterns of proportions between frequencies encountered in all of the dialects of English with t/d deletion.

Using the Euclidean metric one again,[21] the corresponding model frequencies are given in Table 4.16. Table 4.17, represented graphically in Figure 4.3 (note the scale of

TABLE 4.15. Cross-dialectal t/d deletion frequencies

Dialect	#C	#V	##
Philadelphians (Guy 1980)	1.0	0.12	0.38
New Yorkers (Guy 1980)	1.0	0.83	0.66
Chicanos (Santa Ana 1991)	0.61	0.32	0.33
African Americans (Labov et al. 1968)	0.76	0.29	0.73

[21] This time, for each dialect, I solved the problem

$$p^* = argmin\sqrt{(p^v - \#V)^2 + (p^\# - \#\#)^2 + (p^3 - \#C)^2},$$

where the superscript v and # represent the number of violations assigned to that environment according to the dialect's stringency relationships and #X represents the attested retention frequency in the _#X environment for that dialect. In dialects where retention in the _#C environment is 0, CUEC[t d]#>C is not suppressible, so the problem can be reduced to

$$p^* = argmin\sqrt{(p^v - \#V)^2 + (p^\# - \#\#)^2}.$$

TABLE 4.16. Model t/d deletion frequencies

Dialect	p^*	#C	#V	##
Philadelphians (Guy 1980)	0.813	1.0	0.187	0.34
New Yorkers (Guy 1980)	0.366	1.0	0.866	0.634
Chicanos (Santa Ana 1991)	0.704	0.651	0.296	0.296
African Americans (Labov et al. 1968)	0.607	0.776	0.393	0.632

TABLE 4.17. Predicted vs. attested t/d deletion frequencies

Dialect	#C	#V	##
Philadelphians (Guy 1980)	0.0	0.067	0.04
New Yorkers (Guy 1980)	0.0	0.036	0.026
Chicanos (Santa Ana 1991)	0.053	0.024	0.034
African Americans (Labov et al. 1968)	0.016	0.103	0.098

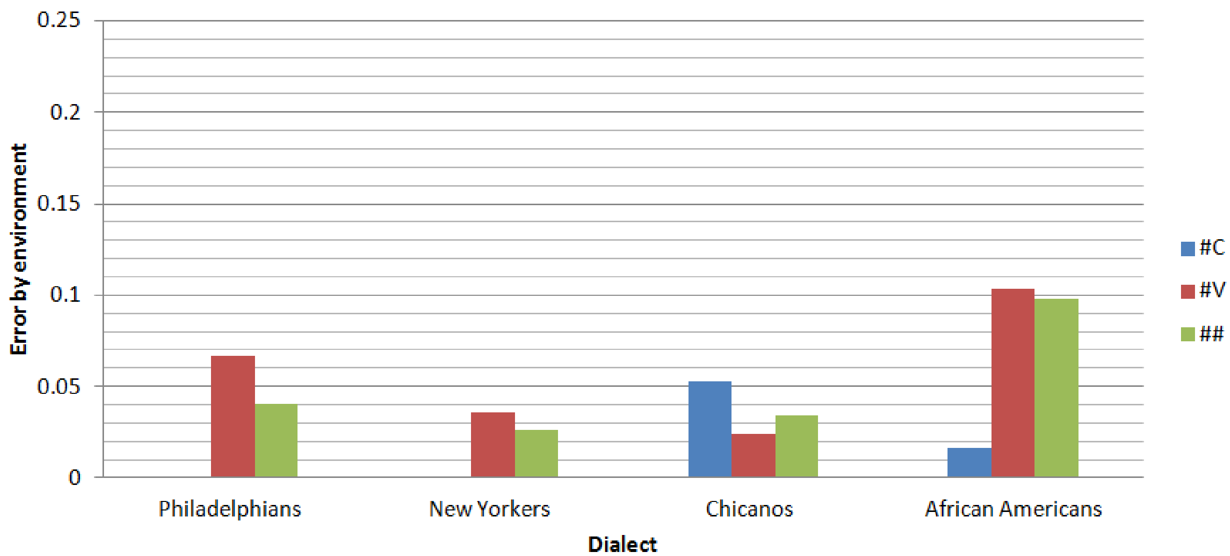

FIGURE 4.3. Predicted vs. attested t/d deletion frequencies (absolute error)

the vertical axis), gives the error in the model's predictions in each environment for these dialects. For the first three dialects, the model frequency is very close to the attested proportions. These dialects represent all of the possible stringency relationships. In Philadelphian English, the _## environment is more marked than the _#V environment; in New York English, the _#V environment is more marked; and in Chicano English, they are on the same level. (CUEC[t d]#>C is not suppressible in the first two.) If these three were the only possible patterns, this analysis would be an unqualified success.

The problem arises in the case of African American English: here, the _# environment is AS MARKED AS the _#C environment, implying that it should receive the same number of violations or, more formally, have the same stringency relationship. Yet our analysis does not permit this to happen because CUEC[t d]#>C can never assign violations to the _# environment given the hierarchy of cues established in (71).

This leaves the _# environment in a position to receive at most two violations from each of the CUEC[t d]# constraints, putting it somewhere in the middle of the two other environments instead of alongside _#C. The solution to this may be as simple as specifying another possible hierarchical relationship between the available cues, namely (79).

(79) Hypothetical hierarchy of cues
_#V > _## = _#C

This could be justified on the basis that formant transitions are preferred to all other cues, and although it is possible to get these in the _#V environment, these cues are not usually available in the other environments.

Alternatively, there is something else going on, and the elaborations on the basic constraint hierarchy in (71)/(78) given in (72) and (79) are off the mark. Depending on the characterization that is made of the data, the empirical description may not translate into a functional Markedness Suppression analysis. The resyllabification analysis given in Kiparsky (1993), for one, has already been shown lacking.

Consider, also, an alternate characterization of the New York dialect, in which the _## environment is preferred to the _#V environment. One possible explanation for this may be that consonants are maintained at the end of a word, but since this is not a CATEGORICAL requirement, a faithfulness constraint cannot be used to enforce it (as faithfulness constraints cannot be suppressible, as would be required). The alternative, FINAL-C (McCarthy 1993), would not work: deletion occurs in the presence of consonant clusters (e.g. *co**st***), meaning FINAL-C will be satisfied whether or not [t d] deletion occurs. The relevant constraint would have to prefer word-final [t] to word-final [s], which is a tough sell, going by Steriade's (1997) account of the cues provided by and available to these segments in those positions.

Ultimately, if there is a problem, it lies in the choice of constraints, and as Coetzee's (2004) analysis makes the same assumptions as I have made here, I will consider this an opportunity for future research on the influence of word-final cues.

In any case, this analysis gives much more accurate model frequencies than the Hungarian analysis above. It is also much better than Kiparsky's (1993): since each dialect only uses a subset of the available patterns, this theory makes no clear predictions about that intradialectal rates of deletion in each environment. Instead, dialects either always adopt one pattern without wavering – say, categorical deletion before C and V – or they vary between no deletion (occurring in three of the six possible rankings between PARSE, ALIGN, and SYLLWF) and one or more of the three remaining patterns (each of which is in one of the six possible rankings) at a rate decided by the number of rankings actually possible. Since there is overlap between rankings that cause deletion in each environment (e.g. two rankings, deletion everywhere and deletion before C and V, cause deletion before V), the deletion rate in a given environment is the proportion of available rankings that cause deletion in that environment. If all six rankings are possible, for example, the rates in Table 4.18 will attain. If, on the other hand, only one of the three rankings that permit deletion is possible, deletion occurs half the time or a third of the time, no matter the environment, since only the ranking of PARSE is relevant to preventing deletion.

Additionally, this model makes the same empirical predictions as Coetzee's (2004) model with respect to the kinds of frequencies available, but it also accounts for the influence of morphological context, which his model does not.

4.4 Conclusion

At the beginning of this chapter, I pointed out two obstacles for a Markedness Suppression analysis of a phenomenon involving interacting environments: first, the difficulty of getting correct frequency predictions and, second, the difficulty of achieving an analysis given the weak ability of suppressible constraints to effectively dominate the constraints they outrank. I have now shown that a Markedness Suppression analysis of

TABLE 4.18. Frequency of deletion under Kiparsky (1993)

Context	Frequency of deletion
a. Before C	3/6
b. Before V	2/6
c. Before #	2/6

Hungarian vowel harmony and English t/d-deletion can, despite the many different environments involved, produce an analysis of these patterns – but these analyses have mixed success modeling attested frequencies for each variant.

By comparing the analyses that failed (or would have failed) – for Hungarian, any analysis where constraints compete against each other to either penalize or reward harmony, and for English, the attempted adaptation of Kiparsky's (1993) analysis – with the ones that worked, the two cases here suggest a method for analyzing similar patterns of environmental interaction in other languages.

In the failed analyses, each of the suppressible constraints did not refer to and "handle" all of the work necessary to distinguish one environment from another; instead, the interaction between those constraints and others was supposed to generate the interaction of the environment.

The attempt to replicate Kiparsky's (1993) analysis of English is especially illustrative. The constraints (Θ) SYLLWF and Θ ALIGN each promote one environment and demote another, and the system relies on the ranking between them to determine which environment is favored. However, the epiphenomenal results of the interaction between these constraints fails to work as intended, leading to the system's failure to correctly represent all of the empirical distributions of frequencies in each environment.

The analyses that DID work used constraints that internalized the work that the failed analyses offloaded onto the ranking system. The result was a set of suppressible constraints, one for each relevant environment, which had NO crucial ranking with respect to one another. This left the problem of modeling the frequencies that result from the environmental interaction. Because the only way Markedness Suppression can manipulate the model frequencies of a particular analysis is to add or subtract violations to certain candidates, the generalization is that one must find a way to give the less frequent candidates more violations, either by using gradient constraints or by having more constraints assign violations to those candidates. In the analyses above, this was accomplished by mapping each environment to a natural hierarchy and using de Lacy's (2003) stringency to dole out different numbers of violations to each output pattern.

Crucially, the analyses that did not work and the analyses that did work relied on different generalizations of the data. For example, Kiparsky's (1993) interpretation of t/d-deletion depicts it as the result of the interplay between constraints penalizing codas and the various ways of fixing them, while the analysis given here relied on Steriade's (1997) theory of licensing by cue. Both describe the same surface phenomenon, but each has a different focus. As the discussion of cue hierarchies in the previous section showed, this is very much a point of contention. Since each of these analyses relied on stringency, barring another way to resolve the problem of assigning more violations to less frequent output forms specific to a given phenomenon, the solution to the interaction problem presented here predicts that differences in output frequencies depending on environments will always pattern according to a natural hierarchy of some kind.

Another question is whether this solution is enough to account for all attested frequency patterns. Once again, the lingering problem interaction presents Markedness Suppression is not whether an analysis is possible at all, but whether that analysis models frequencies correctly. For Hungarian, that was not the case at all. The analysis of English was far closer, but it hinged partly on assumptions made about the phonetics involved.

Still, it is not clear that any other theory does better. The success rate of the

English model was on par with Coetzee's (2004) Rank-Ordered EVAL analysis in its empirical predictions, but that model does not even give numeric frequency estimates. Meanwhile, Stochastic OT, which can model Hungarian almost exactly, does not make the any predictions about possible patterns, allowing for possibly impossible patterns (e.g. the inverse of the height effect). There are, thus, tradeoffs to any theory.

CHAPTER 5

CONCLUSION

5.1 Summary

This thesis began by laying out two potential analytical problems for Markedness Suppression. One, the coordination problem, pertained to the theory's ability to handle patterns in which suppression of a violation of a constraint at one locus of variation appears to be structurally conditioned by decisions made at other loci of variation, in apparent violation of the assumption that the decision to suppress each violation is independent of the decision to suppress any other. The second, the interaction problem, pertained to systems in which variation is conditioned by the phonological environment in which a locus of variation occurs and the interaction between multiple such environments.

For the coordination problem, discussed in Chapter 3, I considered [p]~[b] variation in Warao and vowel reduction in Shimakonde. For Warao, I showed that it is possible to account for the ban on combinations of [p] and [b] in a single word by using consonant harmony to require the same voicing setting for each labial consonant. For Shimakonde, I showed that the "continuity of reduction" pattern, which requires vowel reduction to begin at the left edge of a word and proceed in an unbroken chain from there, can be accounted for with an ALIGN constraint and another constraint blocking "gaps" between reduced vowels. I then compared these two patterns to the pattern in French, in which phonotactics prevent certain schwas from being deleted, concluding that what appears to be coordination between loci is simply the result of an interaction between a variable pattern and one or more other patterns of the language. This also makes the typological prediction that all coordinated variation will follow this pattern: that is, there are not coordinated patterns where the coordination is not due to the structural requirements of some attested pattern (e.g. voicing harmony).

For the interaction problem, discussed in Chapter 4, I considered Hungarian vowel harmony and t/d deletion in English. For both languages, I showed that it was possible to account for their patterns if the relevant constraints referred not to general structural requirements, but to the influence or markedness of each environment that contributes to variation. I also showed that it is possible to ensure almost all of the necessary typological predictions – the height and count effects in Hungarian, in which the height and number of neutral vowels each contribute differently to the rate of suffix harmony, as well as the cross-linguistically constant dispreference for retaining [t d] in the _#C environment in English – by implementing a stringency relationship (de Lacy 2003) between the constraints used in the analysis. This approach should be valid for any systems with interacting environments. However, this approach also makes the typological prediction that interactive patterns are sensitive to natural hierarchies.

5.2 Future research

As summarized above, this thesis advanced two generalized approaches for dealing with the coordination and interaction problems. Both of these are open to further investigation. In particular, the analytical generalizations and typological predictions made on the basis of Warao, Shimakonde, Hungarian, and English can be tested on other languages that also appear to exhibit the coordination and/or interaction patterns.

In addition, most of these analyses had some issues with variant frequencies. For Warao and Shimakonde, a lack of data actually prevented the evaluation of the frequency predictions made by the analyses presented for these languages. For Hungarian, the predicted frequencies were at considerable odds with those observed, although the analysis could correctly account for the broader generalizations, i.e. the direction of the height and count effects. On the other hand, the predictions made for English were fairly accurate.

In the absence of the necessary data on Warao and Shimakonde, it is difficult to see much headway being made there. For Hungarian, there are two ways to go. Currently, the parameter *p* is set on a per-language basis. One possibility is to make Markedness Suppression move closer to Stochastic OT and allow *p* to be set on a per-CONSTRAINT basis instead. As mentioned in Chapter 4, the main tradeoff to doing this is that you lose the ability to restrict the value of *p* so that you must account for the full range of values for *p* for each constraint. This may result in poor predictions when not all combinations of relative frequencies are attested, although this may be mitigated somewhat by things like stringency.

On the other hand, it is also possible to take Coetzee's (2004) critique of frequency estimates to heart and abandon the quest for numerical accuracy altogether. No matter how close numerical estimates of frequencies get to a particular corpus, it may not be worthwhile to be modeling that corpus at all: those aggregate frequencies may not correspond to the grammar of any actual speaker. In this case, the problem of frequency modeling is not a problem at all, as Markedness Suppression does fairly well in describing relative frequencies when the magnitude of those relations is ignored. Both of these possibilities are worth considering.

5.3 Markedness Suppression as a theory of variation

As discussed in Chapter 2, Markedness Suppression is designed as an answer to the harmonic bounding problem of the Partial Orders theory. Whereas Partial Orders allows no room for selective fulfillment of a constraint, Markedness Suppression permits any degree of it for suppressible constraints. This successfully permits intermediate, harmonically bounded candidates to win.

But Markedness Suppression, in permitting so MUCH selectivity in fulfillment of constraints, might have gone the other way. Kaplan (2012) points to Eastern Andalusian vowel harmony, in which harmony is "all or nothing," as a system that the Partial Orders theory inherently models well but which Markedness Suppression cannot handle without using redundant constraints to pare down the set of optimal candidates to those that follow this requirement. This is very similar to my findings for Warao, where Markedness Suppression failed to permit a flat-rate frequency for [p], predicting instead a different rate for each number of loci of variation. As noted above, this could actually be the attested pattern – but if it is not, then Warao is another case of the "all or nothing" problem. In any case, both of the systems could be handled analytically by Markedness Suppression, but neither of them were handled WELL.

In the end, the main obstacle for Markedness Suppression as a theory of variation is not its capacity to develop analyses of a given phenomenon, which was the impetus for investigating coordination and interaction. Instead, the question is whether Markedness Suppression is able to account for these systems in an elegant way. It seems whatever theory of variation we adopt must be able to handle BOTH all-or-nothing systems AND those that favor harmonically bounded candidates. Partial Orders and Markedness Suppression are each designed to reflect only one of these aspects of variation.

As Kaplan (2014) demonstrates, current research in the Partial Orders tradition is making some headway in overcoming these design restrictions. Future research in Markedness Suppression could do the same. Like all other theories of variation, Markedness Suppression does some things better than others and still has plenty of room for improvement. Nothing encountered in any of these analyses is a reason to reject it out of hand; contra Kimper (2011a), it is perfectly functional as a theory of variation.

REFERENCES

ANTTILA, ARTO. 1997. Deriving variation from grammar. *Variation, change and phonological theory*, ed. by Frans Hinskens; Roeland van Hout; and W. Leo Wetzels, 35–68. Amsterdam: Benjamins.

ANTTILA, ARTO. 2007. Variation and optionality. *The Cambridge handbook of phonology*, ed. by Paul de Lacy, 519–536. Cambridge: Cambridge University Press.

BOERSMA, PAUL, and BRUCE HAYES. 2001. Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32.45–86.

COETZEE, ANDRIES W. 2004. *What it means to be a loser : non-optimal candidates in Optimality Theory.* Amherst, MA: University of Massachusetts, Amherst dissertation.

COETZEE, ANDRIES W. 2006. Variation as accessing 'non-optimal' candidates. *Phonology* 23.337–385.

COTE, MARIE-HELENE. 2000. *Consonant cluster phonotactics: a perceptual approach.* Cambridge, MA: MIT dissertation.

CROSSWHITE, KATHERINE. 2001. *Vowel reduction in Optimality Theory*. Rochester, NY: University of Rochester dissertation.

DE LACY, PAUL. 2003. Conflation and hierarchies. *Cambridge Occasional Papers in Linguistics* 1.83–100.

DELL, FRANÇOIS. 1973. *Les règles et les sons: introduction à la phonologie générative.* Paris: Hermann. Translated 1980 by Catherine Cullen as *Generative phonology and French phonology.* Cambridge: Cambridge University Press.

ETTLINGER, MARC. 2009 Input-driven opaque generalizations for harmony and tone. *Lingua,* in revision.

GUY, GREGORY R. 1980. Variation in the group and the individual: the case of final stop deletion. *Locating language in time and space,* ed. by William Labov, 1–36. New York: Academic Press

GUY, GREGORY R. 1991. Explanation in variable lexical phonology: an exponential model of morphological constraints. *Language Variation and Change* 3.1–22.

HANSSON, GUNNAR ÓLAFUR. 2001. *Theoretical and typological issues in consonant harmony*. Berkeley, CA: University of California, Berkeley dissertation.

HAYES, BRUCE, and ZSUZSA LONDE. 2006. Stochastic phonological knowledge: the case of Hungarian vowel harmony. *Phonology* 23.59–104.

ITO, JUNKO, and ARMIN MESTER. 1999a. The structure of the phonological lexicon. *The handbook of Japanese linguistics*, ed. by Tsujimura Natsuko, 62–100. Malden, MA, and Oxford, U.K.: Blackwell Publishers.

ITO, JUNKO, and ARMIN MESTER. (1999b). Realignment. *The prosody-morphology interface*, ed. by René Kager; Harry van der Hulst; and Wim Zonneveld, 188–217. Cambridge: Cambridge University Press.

KAPLAN, AARON. 2008. *Noniterativity is an emergent property of grammar.* Santa Cruz, CA: University of California, Santa Cruz dissertation.

KAPLAN, AARON. 2011. Variation through markedness suppression. *Phonology* 28.331–370.

KAPLAN, AARON. 2012. *Eastern Andalusian vowel harmony and theories of variation.* Salt Lake City, UT: University of Utah, MS.

KAPLAN, AARON. 2014. *Local optionality with partial orders.* Salt Lake City, UT: University of Utah, MS.

KIMPER, WENDELL. 2011a. *Locality and globality in phonological variation.* Pre-publication version. Online: http://personalpages.manchester.ac.uk/staff/wendell.kimper/local_optionality_prepub.pdf

KIMPER, WENDELL 2011b. *Competing triggers: transparency and opacity in vowel harmony.* Amherst, MA: University of Massachusetts, Amherst dissertation.

KIPARSKY, PAUL. 1993. *An OT Perspective on phonological variation.* Handout of talk presented at Rutgers Optimality Workshop I, New Brunswick, NJ.

KIRCHNER, ROBERT. 1996. Synchronic chain shifts in Optimality Theory. *Linguistic Inquiry* 27.341–350.

LABOV, WILLIAM; PAUL COHEN; CLARENCE ROBINS; and JOHN LEWIS. 1968. *A Study of the non-standard English of Negro and Puerto Rican speakers in New York City.* Volume 1. Philadelphia: U.S. Regional Survey.

LIPHOLA, MARCELINO MARTA. 2001. *Aspects of phonology and morphology of Shimakonde*. Columbus, OH: The Ohio State University dissertation.

MCCARTHY, JOHN. 1982. Prosodic structure and expletive infixation. *Language* 58.574–590.

MCCARTHY, JOHN. 1993. A case of surface constraint violation. *Canadian Journal of Linguistics* 38. 69–195.

MCCARTHY, JOHN. 2000. Harmonic Serialism and parallelism. *Proceedings of NELS* 30, ed. by Masako Hirotani, 501–24. Amherst, MA: GLSA Publications.

OSBORN, HENRY. 1966. Warao I: phonology and morphophonemics. *International Journal of American Linguistics* 32.108–123.

PATER, JOE. 1996. *Austronesian nasal substitution and other NC effects.*, University of British Columbia, MS. ROA-160-1196, Rutgers Optimality Archive, http://roa.rutgers.edu.

PENG, LONG. 2000. Nasal harmony in three South American languages. *International Journal of American Linguistics* 66.76–97.

PRINCE, ALAN, and PAUL SMOLENSKY. 1993. *Optimality Theory: constraint interaction in generative grammar.* Rutgers University & University of Colorado, Boulder, MS.

RINGEN, CATHERINE O, and ROBERT M VAGO. 1998. Hungarian vowel harmony in Optimality Theory. *Phonology* 15.393–416.

ROSE, SHARON, and RACHEL WALKER. 2001. *A typology of consonant agreement as correspondence.* NELS 31.

SAMEK-LODOVICI, VIERI, and ALAN PRINCE. 1999. *Optima.* University College London and Rutgers University, MS. ROA-363, Rutgers Optimality Archive, http://roa.rutgers.edu.

SAMEK-LODOVICI, VIERI, and ALAN PRINCE. 2005. *Fundamental properties of harmonic bounding.* University College London and Rutgers University, MS. ROA-785, Rutgers Optimality Archive, http://roa.rutgers.edu.

SANTA ANA, OTTO. 1991. *Phonetic simplification processes in the English of the Barrio: A cross-generational study of the Chicanos of Los Angeles.* Philadelphia, PA: University of Pennsylvania dissertation.

SMOLENSKY, PAUL. 1993. *Harmony, markedness, and phonological activity.* Handout of talk presented at Rutgers Optimality Workshop I, New Brunswick, NJ. ROA-87, Rutgers Optimality Archive, http://roa.rutgers.edu.

STERIADE, DONCA. 1997. *Phonetics in phonology: the case of laryngeal neutralization.* Los Angeles, California: UCLA, MS. Online: http://www.linguistics.ucla.edu/people/steriade/papers/PhoneticsInPhonology.pdf

WALKER, RACHEL. 2000a. *Long-distance consonantal identity effects.* Paper presented at WCCFL 19, Los Angeles, February 2000.

WALKER, RACHEL. 2000b. *Yaka nasal harmony: spreading or segmental correspondence?* Paper presented at BLS 26, Berkeley, February 2000.

WALKER, RACHEL. 2001a. Consonantal correspondence. *Proceedings of the Workshop on the Lexicon in Phonetics and Phonology: papers in Experimental and Theoretical Linguistics,* Vol. 6, ed. by Robert Kirchner, Joe Pater, and Wolf Wikeley, 73–84. Edmonton: University of Alberta.

WALKER, RACHEL. 2001b. Round licensing, harmony, and bisyllabic triggers in Altaic. *Natural language and linguistic theory* 19.827–878.

WALKER, RACHEL. 2011. *Vowel patterns in language.* New York: Cambridge University Press.

VERLUYTEN, PAUL S. 1988. Introduction. *La phonologie du schwa français,* ed. by Paul S. Verluyten, 1–13. Amsterdam: Benjamins.