COMPETING TRIGGERS: TRANSPARENCY AND OPACITY IN VOWEL HARMONY

A Dissertation Presented

by

WENDELL A. KIMPER

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

September 2011

Department of Linguistics

COMPETING TRIGGERS: TRANSPARENCY AND OPACITY IN VOWEL HARMONY

A Dissertation Presented

by

WENDELL A. KIMPER

Approved as to style and content by:

John McCarthy, Chair

Joseph Pater, Member

John Kingston, Member

Caren Rotello, Member

Margaret Speas, Department Chair
Department of Linguistics

ACKNOWLEDGMENTS

A dissertation is never really a solo endeavor, and considerable thanks are due to the many sources of guidance, inspiration, support, and assistance that have contributed to this work.

Thanks are due, first and foremost, to the committee. To my chair, John McCarthy — it has been an immense privilege to work with him over so many years. He has been inspiring, terrifying, encouraging, hilarious, stern, patient, unfailingly wise, and entirely the best advisor one could possibly hope for. His guidance has made me a better phonologist, a better writer of prose, a better instructor, and possibly even a better person. To Joe Pater, who has always challenged me to take seriously ideas and theories that I was initially reluctant to consider (many of which have become integral parts of this dissertation). To John Kingston, who has taught me so much about running experiments, and whose expertise has contributed a valuable perspective to this work. And to Caren Rotello, who put me through my paces in Stats Boot Camp, and inspired a lasting interest in statistical methods.

This work has also benefited from productive discussions with other phonologists — thanks in particular to Gillian Gallagher, Rachel Walker, Elliott Moreton, Peter Jurgec, Maria Gouskova, and Andries Coetzee. Thanks also to Brian Dillon for enthusiastic R help. The Finnish experiment in Chapter 5 was undertaken in collaboration with Riikka Ylitalo (Oulu University). Thanks to Emily Elfner for being the voice of all my auditory stimuli, to Mallory Schleif for running many of the subjects in the discrimination experiment in Chapter 4, and to Clint Hartzell for proofreading assistance (all errors, of course, remain my own). This research was supported by grant BCS-0813829 from the National Science Foundation to the University of Massachusetts Amherst.

My fellow UMass phonology students have been a valuable source of feedback and advice over the years. Many thanks to Kathryn Flack Potts, Matt Wolf, Mike Key, Kat Pruitt, Karen Jesney, Emily Elfner, Brian Smith, Kevin Mullin, Clint Hartzell, Presley Pizzo, Claire Moore-Cantwell, Robert Staubs, and Minta Elsman. Special thanks to Michael Becker, for being an older and wiser academic sibling, and for many cups of tea. My time in grad school has been made more enjoyable and infinitely more survivable by superb officemates Meg Grant, Martin Walkow, Pasha Siraj, and Emily Elfner (and by general fun-inducers Anisa Schardl and Presley Pizzo).

Thanks also to the various other UMass faculty who have offered advice and wisdom (and occasional witty repartee): Ellen Woolford, Kyle Johnson, Lyn Frazier, Seth Cable, Rajesh Bhatt, and Chris Potts. A million thanks to department office staff, Sarah Vega-Liros and Kathy Adamczyk.

Finally, thanks are due to my non-linguist friends and family, who patiently withstood woeful neglect, supported me through moments of panic and doubt, and very kindly refrained from asking me why on *earth* I would want to do a thing like this.

ABSTRACT

COMPETING TRIGGERS: TRANSPARENCY AND OPACITY IN VOWEL HARMONY

SEPTEMBER 2011

WENDELL A. KIMPER

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor John McCarthy

This dissertation takes up the issue of transparency and opacity in vowel harmony — that is, when a segment is unable to undergo a harmony process, will it be skipped over by harmony (transparent) or will it prevent harmony from propagating further (opaque)? I argue that the choice between transparency and opacity is best understood as a competition between potential harmony triggers — segments are opaque when they themselves trigger spreading of the opposing feature value, and transparent when they do not.

The analysis pursued in this dissertation is situated in the framework of Serial Harmonic Grammar, a variant of Optimality Theory which combines the step-wise evaluation of Harmonic Serialism with the weighted constraints of Harmonic Grammar. I argue that harmony is driven by a positively defined constraint, which assigns rewards rather than violations. Preferences for locality and for particular segmental triggers are exerted via *scaling factors* on the harmony constraint — rewards

are diminished for non-local spreading, and increased for spreading from a preferred trigger.

Evidence for this proposal comes from a diverse range of vowel harmony languages, in particular those with multiple non-participating segments which display asymmetries in their amenability to transparency. Segments more likely to be treated as opaque are also independently better triggers — they can be observed to be strong triggers in other contexts, and they are perceptually impoverished along the spreading feature dimension, which means they stand to benefit more from the perceptual advantages conferred by harmony.

This proposal is also supported by experimental evidence. Results of a nonce-word discrimination task and a phoneme recall task both support the claim that harmony is perceptually advantageous; the latter suggests that this advantage obtains even among non-adjacent segments, and I argue that permitting explicitly non-local representations in harmony does not require abandoning phonetic grounding. Evidence for a trigger competition approach comes from a nonce-word study on Finnish disharmonic loanwords, which showed that vowels which are better triggers are more likely to induce transparent harmony, and less likely to be treated as transparent themselves.

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

In languages with vowel harmony, vowels within a morphological or prosodic domain must bear the same value of some feature F. This can be seen both in static properties of stems and in phonological alternations. For example, in Kasem (Gur, Ghana), all vowels within a word agree with respect to their tongue root specification (Casali, 2008).

(1) *ATR harmony in Kasem*

a. sad-ɪ	'grass mat (pl)'	e. kasʊl-ʊ	'sand (pl)'
b. mɪmɪn-ɪ	'thin (pl)'	f. jwaːl-lʊ	'basket (pl)'
c. bəkəd-i	'boy (pl)'	g. gul-lu	'drumming (pl)'
d. kukud-i	'dog (pl)'	h. bwəl-lu	'valley (pl)'

In polysyllabic roots (1b-e), [+ATR] vowels and [–ATR] vowels may not co-occur; each stem contains vowels which agree with respect to their tongue root features. Suffixes alternate based on the [ATR] specification of the root — the plural suffix surfaces as [ɪ,ʊ] when preceded by a [–ATR] stem (1a-b,e-f) and as [i,u] when preceded by a [+ATR] stem (1c-d,g-h).

In Kasem, all vowels are participants in harmony — that is, all vowels are subject to the co-occurrence restrictions in roots, and all vowels alternate in suffixes. Kasem's 10-vowel inventory is fully symmetric with respect to tongue root features — each [+ATR] vowel is paired with a [–ATR] counterpart.

Like Kasem, Yoruba (Niger-Congo, Nigeria) also requires harmony with respect to tongue root features; the words in (2a-b) contain only [+ATR] vowels, while the words in (2c-d) contain only [−ATR] vowels.

(2) *ATR harmony in Yoruba*

a. oko	'farm'	c. ɔkɔ	'husband'
b. ehoro	'rabbit'	d. ɔpɔlɔ	'cheek'

However, Yoruba's 7-vowel inventory is not fully symmetric — the high [+ATR] vowels [i,u] have no [−ATR] counterparts (Archangeli and Pulleyblank, 1994). The high vowels are non-participants (often referred to as neutral) with respect to that language's tongue root harmony system — they do not alternate, and may co-occur in roots with both [+ATR] and [−ATR] vowels. In (3e-h), the [+ATR] high vowel [i] co-occurs with the [−ATR] vowels [ɛ,ɔ].

(3) *Non-participating high vowels in Yoruba*

a. igbe	'noise'	e. ìgbɛ́	'excrement'
b. igbó	'bush'	f. ilɛ̀	'ground'
c. eruku	'dust'	g. itɔ́	'saliva'
d. ìrèké	'sugarcane'	h. ìrɔ̀lɛ́	'evening'

When non-participating vowels co-occur with participating vowels, they may either behave as transparent or as opaque. Transparent vowels may be "skipped over" by harmony — a participating vowel on one side of a transparent vowel may affect a participating vowel on the other side. Opaque vowels, however, prevent further propagation of harmony.

The Ifẹ and Ọyọ dialects of Yoruba illustrate the difference between transparency and opacity. In Ifẹ Yoruba, high vowels may be skipped over by harmony, but in Ọyọ Yoruba, high vowels block propagation of harmony (Pulleyblank, 2004).

(4) *Transparency and opacity in Yoruba*

	Ifẹ	**Ọyọ**	
a.	ɛúrɛ́	eúrɛ́	‘goat’
b.	ɛ̀lùbɔ́	èlùbɔ́	‘yam flour’
c.	ɔ̀títɔ́	òtítɔ́	‘truth’
d.	ɔdídɛ	odídɛ	‘parrot’

The examples in (68) all consist of a neutral vowel flanked by two participating vowels. In Ifẹ, the initial mid vowel takes on the [−ATR] specification of the final vowel, regardless of the intervening [+ATR] high vowel. In Ọyọ, however, that intervening vowel prevents the initial and final vowels from agreeing with each other.

Transparency in harmony has been the topic of considerable discussion in the phonological literature, in part because the data seem to contradict the claim that interactions among phonological elements must be strictly local. In particular, the predominant view has been that assimilation processes may only take place among segments that are adjacent at some level of representation. However, in transparent harmony, non-adjacent vowels assimilate — in (68a), for example, the tongue root features of the initial vowel are determined by the features of the final vowel in Ifẹ, despite the fact that the initial and final vowels are non-adjacent, and the intervening vowel does not participate in harmony.

There have been many attempts to reconcile transparent segments with strict locality. Early literature on harmony proposed that transparent segments were initially participants, but were changed back at later derivational steps (Clements, 1976; Vago, 1976) — this has been adopted in some analyses based in Optimality Theory (OT), by using various mechanisms for derivational opacity (Bakovic, 2000; Bakovic and Wilson, 2000; Walker, 1998). Other proposals referenced Jensen (1974)'s claim that only irrelevant segments may intervene, either by positing that transparent segments were always unspecified along the relevant feature dimension (Kiparsky, 1981; Archangeli and Pulleyblank, 1994) or by proposing that harmony is accomplished via segmental correspondence relations based on featural similarity (Hansson, 2001; Rose and Walker, 2004; Rhodes, 2010; Walker, 2009). Transparent segments have also been re-analyzed as participants, using the fact that they undergo vowel-to-vowel coarticulation as a way of maintaining strictly local interactions (Gafos, 1999; Benus, 2005; Benus and Gafos, 2007).

There has been considerable motivation to preserve strict locality, because it has been commonly believed that this provides a more phonetically grounded account of harmony. However, the proposals which do so are limited in their ability to account for the full range of typological data on transparent segments. In this dissertation, I address the question of strict locality, and I argue (following Cole and Kisseberth 1994; Kaun 1995; Hayes and Londe 2006, and others) that locality is not in fact strict. Furthermore, I argue that strict locality does not provide a more phonetically grounded approach — I present experimental evidence supporting the claim that non-local harmony is still phonetically motivated, and I argue that accounts that attempt to preserve strict locality often end up abandoning phonetic grounding in order to do so. Strict locality has also been argued for on purely formal grounds; I argue further that there is no formal obligation to universally require locality in all phonological interactions.

Of particular interest to the present proposal are languages where some non-participating segments are transparent while others are opaque. In Menominee, for example, height harmony does not affect low vowels (Bloomfield, 1962). The low back vowel [a,a:] is transparent to harmony (5a-b), but the low front vowel [æ,æ:] is opaque (5c-d). In (5a), the initial [e:] in [ce:pa:hkow] becomes [i:] in [neci:pa:hkim], despite the intervening [a:]. In (5c) the initial [o:] in [so:wa:næhki:qsew] is mid, despite a subsequent [i:], because [æ] intervenes.

(5) *Transparency and opacity in Menominee*

a. c**e:**pa:hkow 'he cooks'

nec**i:**pa:hk**i**m 'cook-NOM'

b. s**o:**poma:hkow 'he makes sugar'

nes**u:**poma:hk**i**m 'sugar-maker'

c. s**o:**wa:næhk**i:**qsew 'he has his hair blown back by the wind'

d. p**e:**htæhk**i:**?taw 'he sticks his head in'

Languages like Menominee make it clear that there is no single, language-wide parameter governing locality — non-participating segments differ in the extent to which they are amenable to being skipped. A major focus of this dissertation is to pursue the reasons for this kind of asymmetry — what determines how likely a segment is to behave as transparent?

I argue that the choice between transparency and opacity represents a competition between potential harmony triggers — the local trigger (the non-participant) on the one hand, and the non-local trigger on the other (under this account, opacity is not merely blocking of harmony, but spreading of the opposing feature value). Factors which influence the suitability of each of those competitors as harmony triggers influence the outcome of that choice — segments are more likely to be transparent when they are poor triggers, and more likely to be opaque (initiating their own harmony domain) when they are good triggers. Likewise, transparent harmony is more likely to be initiated from a segment which is a good trigger, and less likely from a segment which is a poor trigger.

This dissertation focuses specifically on vowel harmony processes; set aside are issues of consonantal participation in harmony (see e.g. Gafos 1999; Ni Chiosain and Padgett 2001), harmony systems in which primarily consonantal features spread (see e.g. Hansson 2001; Rose and Walker 2004) or harmony systems in which consonants and vowels interact, such as nasal harmony (see e.g. Walker 1998). I will assume henceforth that consonants are not participants in vowel harmony processes, but this is an assumption adopted primarily for purposes of scope and exposition, and the question of consonantal participation is one which merits further exploration.

The thesis is organized as follows. In Chapter 2, I lay out the basic framework of an account of harmony in general, and an account of non-participation in particular. I propose that harmony is motivated by a positively-formulated constraint, which assigns rewards for spreading rather than penalties for disagreement. I show how adopting the framework of Harmonic Serialism (McCarthy, 2000, 2007a,b), with its gradually harmonically improving derivations, permits the inclusion of a positively formulated constraint (which would encounter the problem of infinite goodness in versions of OT with unrestrained candidate sets). I compare the proposed positive constraint with previous approaches to formulating negative harmony constraints, arguing that a positive constraint situated in Harmonic Serialism provides a more typologically sound theory of harmony than previous approaches.

Chapter 2 also discusses the range of factors that can cause a segment to fail to participate in harmony. I argue that feature co-occurrence restrictions, while important in accounting for a majority of cases of non-participation (and capturing the generalization that non-participation frequently accompanies inventory gaps), are neither necessary nor sufficient to prevent harmony. Segments may undergo harmony despite violations of otherwise-obeyed markedness constraints, either engaging in non-contrastive alternations or re-pairing by changing along another feature dimension. Furthermore, segments whose harmonic counterparts are present in a language's inventory can still nonetheless fail to undergo harmony when they appear in exceptional morphemes or lexical strata, or when they are protected by perceptually-based asymmetries in faithfulness.

Chapter 3 addresses the factors influencing the choice between transparency and opacity. I argue that this choice is best understood as a competition between potential harmony triggers — the non-local segment bearing the dominant feature on the one hand, and the local segment bearing the non-dominant feature on the other. Overall preferences for dominance or locality can influence the choice across the board, but segment-specific factors which influence the strength of each of the triggers can interact with those preferences. I show that in each of the languages where non-participants pattern differently — with some non-participating segments treated as transparent while others are opaque — the asymmetry can be explained by differences in the suitability of each of the segments as harmony triggers. Segments which are better triggers (by virtue of being perceptually impoverished, and therefore in most need of the advantages conferred by harmony) are more likely to be opaque, and segments which are poor triggers are more likely to be treated as transparent.

I propose that this should be formally represented via a scaling factor on the constraint motivating harmony. Rewards are decreased for spreading when the trigger and target are non-local, and rewards are increased when the trigger is comparatively perceptually impoverished. Cumulative constraint interaction in Harmonic Grammar (Smolensky and Legendre, 2006; Pater, 2009; Potts et al., 2010, and others) enables the difference in rewards earned for spreading under different triggering conditions to influence the competition between transparent and opaque harmony. Furthermore, I propose that the scaling factors are defined operationally rather than representationally, capturing the generalization that conditions on the locality or quality of harmony triggers only ever exert their influence by preventing harmony (and not by inducing other changes).

Chapter 4 takes up the question of strict locality, and the phonetic motivations for harmony. I discuss the evidence for and against treating transparent segments as undergoers of harmony, concluding that they are in fact skipped, and I argue that a theory of harmony based on strict articulatory locality will fail to account for transparent harmony without also losing phonetic grounding in the process. I present experimental evidence that harmony is perceptually beneficial — listeners are faster and more accurate at identifying vowel contrasts in harmonic contexts than in disharmonic contexts. Furthermore, I present experimental evidence that this benefit is obtained even in the context of harmony among non-adjacent segments. Based on this evidence, I argue that strict locality does not in fact provide a more phonetically grounded theory of harmony than one which includes explicitly non-local interactions.

In Chapter 5, I test the prediction that segments which are preferred triggers will not only be more likely to be treated as opaque, but they will also be more likely to induce transparent harmony. This chapter presents the results of a nonce-word study on Finnish loanwords, where toleration of disharmonic stems leads to the non-participation of segments which would otherwise undergo harmony. When these exceptional non-participants intervene between a vowel bearing the dominant feature and the target in the suffix, their behavior is variable — they may either be transparent or opaque. In this experiment, subjects were more likely to choose transparency when the non-participant was itself a poor trigger, but also when the non-local trigger was a good trigger. I discuss the problems these results pose for other theories of transparency, and show that the competing triggers approach alone captures the full range of experimental results.

Finally, Chapter 6 compares the proposal advanced in this dissertation with other extant theories of transparency and opacity. I argue that theories which maintain strict locality struggle to account for the behavior of transparent segments (and the conditions under which transparency is possible). I also address the formal arguments in favor of representing locality as a universal, inviolable restriction, concluding that there is no *a priori* reason to prefer a strictly local analysis, and that the present approach provides better empirical coverage. I also discuss other proposals which relax the requirement of locality, arguing that the predictions of the competing triggers approach are a better fit with the typology of harmony processes than the alternatives.

CHAPTER 2

HARMONY AND NON-PARTICIPATION

2.1 Introduction

A discussion of vowel harmony from the perspective of an Optimality Theory (OT) based analysis must necessarily begin by defining the constraint which drives the process. The nature of such a constraint, and its formal definition, has been the subject of considerable debate in the literature. Broadly speaking, common approaches include agreement constraints, which assign violations to adjacent segments with different feature values (Bakovic, 2000; Eisner, 1999; Lombardi, 1999, 2001; Pulleyblank, 2004), alignment constraints, which require that the edges of feature domains align themselves with the edges of prosodic or morphological constituents (Archangeli and Pulleyblank, 2002; Cole and Kisseberth, 1994, 1995a; Kirchner, 1993; Pulleyblank, 1996; Smolensky, 1993), and Agreement by Correspondence (ABC), which requires segments to be in surface correspondence relations, and requires that corresponding segments agree (Hansson, 2001; Rose and Walker, 2004; Walker, 2009; Rhodes, 2010).

While these proposals have been successfully deployed in many analyses of harmony processes, they encounter difficulties when the range of their typological predictions is taken into consideration. In this chapter, I argue for an alternative way of driving harmony — a positively formulated constraint which assigns rewards for spreading (rather than assigning penalties for disagreement). I show how the proposed constraint, situated in the framework of Harmonic Serialism (HS; a derivational variant of OT), avoids the problematic predictions of previous, negatively defined constraints.

In addition to presenting the basic mechanism for driving harmony, this chapter also discusses a variety of factors which can prevent harmony from occurring. In particular, I discuss how feature co-occurrence restrictions, lexical exceptions, and faithfulness asymmetries can cause segments to be non-participants, failing to undergo harmony — the treatment of these non-participating segments is primary focus of this dissertation, and is pursued in subsequent chapters.

The chapter is organized as follows. The remainder of this chapter provides a background introduction to Harmonic Serialism. Section 2.2 introduces and defines the positively formulated harmony constraint; Section 2.2.1 discusses how Harmonic Serialism enables positive constraints to be a viable possibility (by avoiding the infinite goodness problem encountered by positive constraints in parallel versions of OT), while Section 2.2.2 discusses the problematic predictions made by negatively defined agreement and alignment constraints, and shows how a positive constraint is able to avoid those problems. Section 2.2.3 discusses the similarities and differences between the present approach and OT with Targeted Constraints (TCOT; Wilson 2004, 2006). Section 2.3 discusses some further predictions of this overall approach to harmony. Finally, Section 2.4 discusses factors which can interfere with harmony, showing how interaction with constraints against marked feature combinations (Section 2.4.1) or specific types of faithfulness constraints (Sections 2.4.2 and 2.4.3) can cause segments to fail to undergo harmony.

2.1.1 Background: Harmonic Serialism

Harmonic Serialism (HS) is a variant of OT; it shares much of its fundamental architecture for constraint interaction and evaluation, but with a restricted GEN and multi-step derivation that involves repeated passes through EVAL (McCarthy, 2000, 2007a,b).

In HS, Gen is restricted to producing candidates that each differ from the input by at most a single application of a single change. What constitutes a single change is a matter of ongoing research, and is determined empirically rather than *a priori*; see for example McCarthy (2010) for a discussion of the kinds of evidence brought to bear on the question.

The limited candidate set generated by a restricted Gen is evaluated by Eval, just as in parallel versions of OT. However, instead of exiting the derivation as the surface form, the optimum is sent back to Gen to serve as the input to a new step in the derivation. Once again, candidates are generated which differ from this new input by at most a single instance of a single change, and this candidate set is sent to Eval for evaluation. The Gen→Eval loop continues until the changes made by Gen no longer increase harmony — the derivation converges when the faithful candidate at the current step is chosen as optimal. The difference between the parallel evaluation of OT and the serial evaluation of HS is visualized in (6).

(6) a. **Parallel Evaluation**

$$\begin{tikzcd} \text{Input} \arrow[r] & \boxed{\text{Gen}} \arrow[r] & \boxed{\text{Eval}} \arrow[r] & \text{Output} \end{tikzcd}$$

b. **Serial Evaluation**

$$\begin{tikzcd} \text{Input} \arrow[r] & \boxed{\text{Gen}} \arrow[r] & \boxed{\text{Eval}} \arrow[r] \arrow[l, bend left=60] & \text{Output} \end{tikzcd}$$

As a consequence of this way of proceeding, HS derivations are necessarily gradual and harmonically improving. Because Gen is restricted, each step in the derivation can advance by at most a single change (gradualness). Because each step involves choosing an optimum at Eval, and the derivation converges if the faithful candidate is chosen, each step must improve performance with respect to the constraint hierarchy (harmonic improvement).

The gradual harmonic improvement of HS has been shown to have a number of typological advantages; see for example Pruitt (2008) on foot parsing, Pater (2008c) on syllabification, Jesney (2009) on positional faithfulness, Elfner (2009) on stress/epenthesis interactions, McCarthy (2008c) on cluster simplification, McCarthy (2008b) on metrically conditioned syncope, and McCarthy (2007a); Wolf (2008) on opacity (both phonology-phonology and phonology-morphology).[1] In particular, McCarthy (to appear) for a discussion of the advantages of HS in analyzing vowel harmony processes.

2.2 A Positive Harmony Imperative

In OT-based frameworks, constraints have traditionally been negatively defined (assigning violations) rather than positively defined (assigning rewards). This means that harmony-driving constraints have been designed to punish disharmony. Instead, I propose a positively defined constraint that specifically rewards assimilation, defined in (7).

(7) SPREAD(±F): For a feature F, assign +1 for each segment linked to F as a dependent.

The definition of this constraint relies on fairly standard assumptions about autosegmental representations — each feature is associated with at most one head, and additional segments associated with that feature are dependents. I set aside, for now, the question of which segment is awarded head status; I assume (following Smolensky 2006, and others) that the segment which bears F underlyingly is the head, but

[1]The analysis of opacity (McCarthy, 2007a; Wolf, 2008) is situated in Optimality Theory with Candidate Chains (OT-CC), a serial version of OT which shares many properties with the version of HS discussed in this dissertation in other recent work, but with some significant formal differences.

nothing crucial rests on this assumption (and other mechanisms of determining head status would be perfectly compatible with the analysis presented here).

Because HS restricts Gen to producing candidates which differ from the input by at most a single instance of a single change, it is important to identify the single change involved in harmony. In other words, what is the operation that Gen performs when harmony takes place? In this case, the relevant operation is the creation of an autosegmental association between a feature and a segment; this is illustrated in (8a).

If a segment is already associated with another instance of ±F, that association may be undone as part of the linking operation.[2] In other words, de-linking comes for free as part of creating an autosegmental association. This is illustrated in (8b).

(8) For a feature $\pm F_i$ and a segment S, create an autosegmental association between F_i and S. If S is associated with another instance $\pm F_j$, the link between S and $\pm F_j$ is dissolved.

a. Autosegmental Linking:

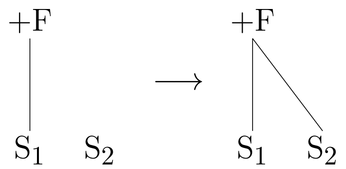

b. Autosegmental Linking (with de-linking)

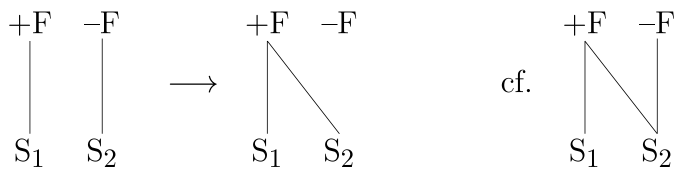

There are at least two faithfulness constraints, defined in (9),[3] violated by the linking operation above. The constraint Dep(link) is always violated whenever the

[2] I leave aside the question of whether de-linking is *always* involved; representations involving a single segment linked to multiple autosegments have long been used to represent contour tones, and so should not be excluded as possible representations. However, nothing parallel to contour tones seems to arise from vowel harmony processes.

[3] For constraints other than Spread(F), I retain the traditional negative formulation.

linking operation is performed, because it penalizes the addition of a new autosegmental association. On the other hand, IDENT(F) is not *necessarily* violated by an instance of linking — it is only violated when the instance of a feature a segment becomes associated with has the opposite value from the feature it was previously associated with. The distinction between DEP(LINK) and IDENT(F) will be discussed further in Sections 3.3.3 and 3.4.1; for most instances of harmony, they are in a stringency relationship (whenever IDENT(F) is violated as a result of harmony, DEP(LINK) is violated, but not vice versa), so the more specific IDENT(F) will be used wherever the distinction is not relevant.[4]

(9) a. DEP(LINK): Assign –1 for an autosegmental association in the output that is not present in the input.

b. IDENT(F): Assign –1 for a segment that is αF in the input, but $-\alpha$F in the output.

When SPREAD($\pm$F) dominates faithfulness, harmony is the result. In Kasem, for example, there is a root-controlled process of tongue root harmony; affixes alternate to take on the ATR features of the root (Casali, 2008). This is illustrated in (10) with the Group C singular suffix, which is [-ʌ] with [+ATR] roots, and [-a] with [–ATR] roots.

(10) **Kasem Group C Singular: [-ʌ]~[-a]**

[4]It is also possible that a third faithfulness constraint, MAX(LINK), is involved in instances where autosegmental linking requires delinking of a previously associated feature. In practial terms, the inclusion of such a constraint will have an effect only in cases of constrastive underspecification — with target segments specified for either feature value, linking will entail delinking. The issue of contrastive underspecification is set aside here, but deserves further consideration.

[−atr] Roots		[+atr] Roots	
dzɪn-a	‘hand’	dig-ʌ	‘room’
nag-a	‘leg’	lʌŋ-ʌ	‘song’
zʊn-a	‘calabash’	bug-ʌ	‘river’
kɔg-a	‘back’	tʃoŋ-ʌ	‘path’

If the input contains a root with a [+atr] vowel, a ranking of Spread(+atr) ≫ Ident(ATR) ensures that the suffix vowel will always be [+atr] regardless of its input specification. This is demonstrated in the tableau in (11); the faithful candidate (11a) earns no rewards on either Spread(+atr) or Spread(−atr), and incurs no violations of Ident(ATR). In candidate (11b), an autosegmental association has been created between the [+atr] feature and the suffix vowel. This earns a reward on Spread(+atr), at the cost of violating Ident(ATR); because Spread(+atr) ≫ Ident(ATR), the candidate with harmony (11b) is the winner.[5] At the next step of the derivation, no further improvement is possible — there are no additional vowels to link to — so the derivation will converge.[6]

(11) **Step 1**

[+] [−] dig-a	Spread(+atr)	Spread(−atr)	Ident(ATR)
a. [+] [−] dig-a	W		L
b. ☞ [+] [−] dig-ʌ	+1		−1

Step 2: *Convergence*

[5]There is a notable omission from the candidate set considered here, namely a candidate where [−atr] has spread from the suffix vowel to the root vowel. I assume that root control is accomplished via positional faithfulness (Beckman, 1998) — a ranking of Ident(ATR)$_{\text{Rt}}$ ≫ Ident(ATR)$_{\text{Aff}}$ will ensure that harmony that targets affixes is consistently preferred over harmony that targets roots.

[6]It is important to note that convergence is always with respect to some particular constraint set; given this set of constraints, no further improvement is possible, and the derivation will converge at Step 2.

In Kasem, both [+ATR] and [–ATR] spread — that is, the root always determines the tongue root features of affix vowels, regardless of whether that root is [+ATR] or [–ATR]. This means that SPREAD(–ATR) also dominates IDENT(ATR); the tableau in (12) is similar to the tableau in (11) above, but this time the input contains a root with a [–ATR] vowel. The faithful candidate (12a) earns no rewards on SPREAD(–ATR), but incurs no violations of IDENT(ATR). The winning candidate (12b) has spread [–ATR] to the suffix vowel; it earns a reward on SPREAD(–ATR), at the cost of violating the lower-ranked IDENT(ATR). As above, no further improvement on this constraint hierarchy is possible, and the derivation converges at Step 2.

(12) **Step 1**

[–] [+] zʊn-a	SPREAD(+ATR)	SPREAD(–ATR)	IDENT(ATR)
a. [–] [+] zʊn-a		W	L
b. ☞ [–] [+] zʊn-a		+1	–1

Step 2: *Convergence*

In Kasem, it's not possible to determine the relative ranking of SPREAD(+ATR) and SPREAD(–ATR) — both dominate faithfulness, and harmony is consistently root-controlled. In other words, neither [+ATR] nor [–ATR] can be said to be the dominant feature value. This contrasts with languages like Nez Perce (Hall and Hall, 1980), where a [–ATR] vowel anywhere in the word triggers harmony. The evidence needed to ascertain which feature value is dominant is often difficult to come by; see for example Casali (2008) for a discussion. The role that feature dominance plays in the present theory will be addressed further in Chapter 3.

2.2.1 Infinite Goodness and the Serial Solution

Positive constraints have not been a viable option in OT, because they encounter the Infinite Goodness problem. For any structure that is rewarded, there can be candidates with infinitely increasingly many instances of that structure, and so there ceases to be an optimum (Prince, 2007). This arises because positive and negative constraints are not simply the mirror image of each other — they assign marks differently, and exert different preferences. Compare, for example, the positively-defined SPREAD(+ATR) above and a common negatively defined harmony constraint like ALIGN(+ATR) (which requires that the edge of the [+ATR] domain correspond to the edge of the word). Considering again the Kasem example above, we can compare the preferences of the two constraints in the tableau in (13).

Both constraints prefer a candidate with harmony (13b) over the faithful candidate (13a). However, they differ on candidate (13c), which has inserted an epenthetic vowel and spread [+ATR] to it. The negatively-defined ALIGN(+ATR) does not distinguish candidates (13b) and (13c) — both show equally good alignment — and as a result, (13c) will be (correctly) harmonically bounded; it has not improved performance on the harmony constraint and has gratuitously violated faithfulness. The positive SPREAD(+ATR), however, does distinguish (13b) and (13c) — because of the epenthetic vowel, (13c) has another dependent, and performs better on the harmony constraint.

The problem with this is immediately apparent. If one epenthetic dependent vowel is good, two will be even better, and three will be even better than that — and so on. Because GEN is unrestrained, for any number of epenthetic harmonized vowels, there exists a candidate with one more. There is no longer any best candidate, any single optimum.

(13)

[+] [−] dig-a	SPREAD(+ATR)	ALIGN(+ATR)
a. [+] [−] dig-a		−1
b. [+] [−] dig-ʌ	+1	
c. [+] [−] dig-ʌʌ	+2	
d. [+] [−] dig-ʌʌ+$\aleph_0$	+∞	

The Infinite Goodness problem means that, in parallel versions of OT, any positive constraint is simply untenable. This is true regardless of the particular constraints and their formulation — if a constraint rewards a particular output structure, the candidate set can contain candidates with infinitely many instances of that structure. In HS, however, the gradual harmonic improvement that results from a restricted GEN enables positive constraints to avoid the Infinite Goodness problem. This solution is dependent on the formulation of the positive constraints and on assumptions about the set of possible operations (see Section 2.3.2 for futher discussion) but is fairly robust across a variety of constraint types and fairly reasonable assumptions about gradualness.

The problematic candidates in (13cd) differ from the input by more than a single instance of a single change. Candidate (13c) differs from the input not just in the addition of an epenthetic vowel, but also in the addition of an autosegmental link. It is therefore not a member of an HS candidate set; and, by extension, neither are candidates with additional epenthetic linked vowels.

The tableau in (14) shows how an HS derivation avoids Infinite Goodness. At the first step, as above, there is a faithful candidate (14.1a) and a candidate with harmony (14.1b). There is also a candidate with an epenthetic vowel (14.1c), but crucially there is no candidate like (13c) — no candidate with both epenthesis and harmony. Because Spread(+atr) outranks Ident(+atr), the candidate with harmony wins. This is then the input to the second derivational step. The faithful candidate (14.2a) already has every vowel associated with [+atr], and continues to enjoy its reward for the dependent suffix vowel. Again we can consider a candidate with epenthesis (14.2b), but again this candidate may have *only* epenthesis. Because epenthesis and autosegmental association are two independent operations, they cannot occur simultaneously, and there is no candidate with both epenthesis and harmony. The epenthetic candidate gratuitously violates faithfulness without gaining any further satisfaction from Spread(+atr), and is harmonically bounded; the faithful candidate is optimal, and the derivation converges.

(14) **Step 1**

[+] [−] dig-a	Spread(+atr)	Spread(−atr)	Ident(ATR)	Dep
a. [+] [−] dig-a	W		L	
b. ☞ [+] [−] dig-ʌ	+1		–1	
c. [+] [−] dig-aʌ	W		L	W–1

Step 2: *Convergence*

[+] [−] dig-ʌ	Spread(+atr)	Spread(−atr)	Ident(ATR)	Dep
a. ☞ [+] [−] dig-ʌ	+1			
b. [+] [−] dig-ʌʌ	+1			W–1

It is not possible to epenthesize for the purpose of increasing satisfaction of Spread(F) because the gradual, harmonically improving derivations are myopic.[7] That is, there is no derivational look-ahead — no way of knowing that epenthesis at

[7]See Wilson (2004) for the first use of this term.

Step 2, if chosen, could lead to additional rewards for spreading at Step 3. Because candidates may differ from the input only by a single instance of a single change, and because single changes which are not in and of themselves harmonically improving do not survive to future derivational steps, there is no way to infinitely continue to satisfy a constraint like SPREAD(F) in HS. The Infinite Goodness problem is avoided, and positive constraints are a viable possibility.

It is also not possible to satisfy a constraint like SPREAD(F) with spontaneous feature changes. The first reason for this is that SPREAD(F) rewards only *dependent* segments rather than any segment associated with F — no further reward is earned for e.g. simply changing a [–ATR] segment into a [+ATR] segment. This is illustrated in (15); the input contains two [-ATR] vowels, and candidate (15a) remains faithful while candidate (15b) changes one of those vowels to [+ATR]. Neither candidate obtains any advantage with respect to SPREAD(±ATR), and the unfaithful candidate is elimiated by IDENT(ATR).

(15) **Step 1**

[–] [–] dɪg-a	SPREAD(+ATR)	SPREAD(–ATR)	IDENT(ATR)	DEP
a. ☞ [–] [–] dɪg-a				
b. [+] [–] dig-a			W–1	

Additionally, gradualness in HS means that it is not possible to simultaneously change the value of a segment and spread that feature value; in (15) there is no candidate [digʌ] (which would gain advantage on SPREAD(+ATR)).

Because avoiding the Infinite Goodness problem relies crucially on the gradualness of HS derivations, the overall success of positive constraints depends on what precisely GEN is capable of. For SPREAD(F), and the autosegmental linking operation defined above, derivations will be sufficiently gradual to prevent infinitely increasing satisfaction. But the viability of other positive constraints will depend on the operations that satisfy them, and further exploration of the question of what counts as a single change. For further discussion of the generality of this solution to the Infinite Goodness problem, see Section 2.3.2.

2.2.2 The Pathologies of Negativity

The preceding section argued that a positive spreading constraint is possible; this section argues that such a constraint is also desirable. In particular, I discuss the problematic typological predictions that constraints of the agreement and alignment types make.[8] Wilson (2004, 2006) points out that commonly-used harmony constraints make problematic predictions in systems where spreading is blocked; the discussion of these predictions centers around Malay, where nasal harmony is blocked by obstruents (Onn, 1980; Walker, 2000) — the relevant data are given in (16).

(16) **Nasal harmony in Malay**

a. Unbounded rightward spreading

mĩnõm	'to drink'
bãŋõn	'to rise'
mãɟ̃ãn	'stalk (palm)'
pəŋə̃ŋãh̃ãn	'central focus'
mə̃nãw̃ãn	'to capture'

b. Blocked by obstruents

mãkan	'to eat'
pəŋãw̃ãsan	'supervision'

Constraints of the agreement type, which penalize $[\alpha F][-\alpha F]$ sequences, encounter the Sour Grapes problem (Wilson, 2004; McCarthy, 2004, to appear) — if complete agreement is impossible, no spreading is predicted to occur at all. Take, for example, an input like /pəŋawasan/ — a constraint like AGREE(NAS) will assign a

[8] For a discussion of the Agreement by Correspondence approach to harmony, see Chapter 6.

violation for the [+NAS][–NAS] sequence [ŋa]. However, the attested output form [pəŋã w̃ ãsan], where nasality has spread several segments to the right, performs no better — AGREE(NAS) will assign a violation to the [ãs] sequence. Harmony has not improved performance on AGREE(NAS), and has gratuitously violated faithfulness, so it is predicted to be harmonically bounded. Nevertheless, languages like Malay, rather than giving up whenever total agreement is impossible, spread as far as possible.

SPREAD(F) avoids the Sour Grapes problem because it rewards harmony rather than penalizing disagreement. Because a reward is earned for each segment that harmonizes, spreading iteratively is preferred over failing to spread, even when total agreement is impossible. Indeed, this is something of a prerequisite for harmony constraints in HS in general — because derivations are myopic, no step where spreading occurs has access to information about the viability of continuing to spread.

Alignment constraints, which assign violations for segments intervening between the edge of a feature domain and the edge of a designated morphological or phonological constituent, do not suffer from the Sour Grapes problem. However, they do suffer from a different set of pathologies — because they assign violations to segments that have not harmonized, they can interfere with processes that affect the number of segments in a word (Wilson, 2004, 2006).

For example, in words where harmony is blocked, a sufficiently highly ranked ALIGN can block epenthesis. Consider Malay′, where ALIGN dominates a constraint against word-final consonant clusters (which in turn dominates DEP). When harmony is blocked, an epenthesized segment increases the distance between the edge of the feature domain and the word edge — ALIGN will prefer a candidate with no epenthesis (17a) over a candidate with an epenthetic segment (17b).

(17)

nawakast	ALIGN-R(NAS,PWD)	*CC#	DEP
a. ☞ nãw̃ãkast	–4	–1	
b. nãw̃ãkasət	W–5	L	W–1

In Malay′, final clusters are normally resolved by epenthesis (18a). However, in words where spreading is blocked, no epenthesis occurs (18b). More broadly, a constraint like ALIGN-R can produce a language where harmony does not occur, but epenthesis is blocked to the right of a nasal segment. These patterns are unattested, in Malay or elsewhere (Wilson, 2004, 2006).

(18)

	ALIGN-R(NAS,PWD)	*CC#	DEP
a. tawakasət ∼ tawakast		W	L
b. nãw̃ãkast ∼ nãw̃ãkasət	W	L	W

ALIGN blocks epenthesis because it assigns a violation to each segment intervening between the edge of the feature domain and the edge of the word; epenthesizing a segment, then, adds to the number of violations assigned by ALIGN. Broadly speaking, the problem is that ALIGN is sensitive to the number of non-harmonized segments in a word; a number of the other pathologies identified by Wilson (2004) (including those related to deletion, allomorph selection, and reduplication) also arise because of this kind of segment-counting.

McCarthy (to appear) notes that serial evaluation plays a role in resolving some of the problematic predictions of ALIGN and other harmony constraints (including, to some extent, the Sour Grapes problem). However, sensitivity to the number of non-harmonized segments in a word remains — segment-counting pathologies like harmony-dependent epenthesis blocking are still predicted with ALIGN in HS. The competitors at the nth step of an HS derivation in (19) are exactly the same as those in the OT derivation in (17).

(19) **Step *n***

nãw̃ãkast	ALIGN-R(NAS,PWD)	*CC#	DEP
a. ☞ nãw̃ãkast	–4	–1	
b. nãw̃ãkasət	W–5	L	W–1

Because a positively-defined spreading constraint assigns rewards for autosegmental associations rather than assigning violations for failure to assimilate, it will not be sensitive to the number of non-harmonized segments present. Unlike Align, a positive spreading constraint will not suffer from segment-counting pathologies.[9]

At the first step of the derivation in (20), the candidate with nasal spreading wins over the faithful candidate. This is again the case at the second and third steps; spreading will continue to iterate, segment by segment, until there are no more segments left unassociated or until spreading is blocked. The latter is true in (20); at the fourth step, spreading conflicts with *NasObs. The faithful candidate is chosen as optimal, and the derivation converges.

(20) **Step 1**

nawakast	*NasObs	Spread(nas)	Ident(nas)
a. nawakast		W	L
b. ☞ nãwakast		+1	-1

Step 2

nãwakast	*NasObs	Spread(nas)	Ident(nas)
a. nãwakast		W+1	L
b. ☞ nãw̃akast		+2	-1

Step 3

nãw̃akast	*NasObs	Spread(nas)	Ident(nas)
a. nãw̃akast		W+2	L
b. ☞ nãw̃ãkast		+3	-1

Step 4: *Convergence*

nãw̃ãkast	*NasObs	Spread(nas)	Ident(nas)
a. ☞ nãw̃ãkast		+3	
b. nãw̃ãk̃ast	W-1	L+4	W-1

Because it assigns rewards for harmony rather than violations for failure to harmonize, Spread(nas) will not be sensitive to the number of segments to the right of the blocking obstruent in [nãw̃ãkast]. Consider a candidate with epenthesis to

[9]See McCarthy to appear for another proposal to address this class of pathologies.

break up the final cluster — both the faithful candidate (21a) and the candidate with epenthesis (21b) perform equally well on Spread(nas).

(21) **Step 4**

nãw̃ãkast	Spread(nas)	*CC#
a. nãw̃ãkast	+3	W-1
b. ☞ nãw̃ãkasət	+3	

A positive spreading constraint, therefore, cannot block epenthesis. Furthermore, a constraint like Spread(F) is prevented from inducing epenthesis by the restricted Gen in HS, as we saw in the discussion in the previous chapter. While it is not sensitive to the number of non-harmonized segments in a word, Spread(F) is sensitive to the number of segments associated with a given feature; see Section 2.3.1 for a discussion of the predictions that this makes.

While the problematic predictions of constraints like Agree and Align emerge most strikingly under conditions where harmony is blocked, it should be noted that this is not a consequence of extending those constraints to account for patterns of locality; these problems emerge because of fundamental aspects of how these constraints assign marks. The present proposal, then, is not an extension or refinement on these existing constraints, but rather a basic departure from previous mechanisms driving harmony.

2.2.3 Comparison with Targeted Constraints

The proposal outlined above is somewhat similar to (and, indeed, inspired by) the Targeted Constraints (TCOT) proposal in Wilson (2004, 2006, to appear). In TCOT, markedness constraints evaluate both the marked configuration and the repair used to eliminate it — penalties are assigned to the former, and rewards (or cancellation of violations) are assigned to the latter.

A schematic version of Wilson (2004)'s harmony constraint is given in (22). The markedness component is represented by the parameter λ, and is similar to an Agree constraint — it penalizes adjancent segments which disagree with respect to F. The specified repair is represented by the parameter δ — this constraint can only be satisfied by a change in $\pm$F.

(22) T:SPREAD-L,R([αF],D)

λ A non-[αF] segment immediately to the {left,right} of an [αF] segment in the same domain D.

δ [0F]/[$-\alpha$F]$\rightarrow$[αF]

The specific mechanisms of evaluation have been defined somewhat differently over time; for the purposes of illustration, I will abstract away somewhat from the finer details of computational implementation and instead focus on the overall thrust of the proposal, summarizing the latest instantiation in Wilson (to appear), which is instantiated within the HS framework. Marks for targeted constraints are assessed differently from traditional OT markedness constraints. For each input–output pair, if either the input or the output contain an instance of the marked structure λ, a violation is assessed (and if neither the input nor the output contain an instance of λ, no violations are assessed). However, If the input contains an instance of λ which is repaired in the output via the specified mapping δ, no violations are assessed for the input violation, but instead a reward is assessed for exhibiting the preferred repair.

So, considering the constraint in (22) above; for an input containing a disharmonic sequence, a faithful candidate derived from it will earn a violation (–1), a candidate derived from that input which resolves the disharmonic sequence by e.g. deletion will still receive a violation (–1), but a candidate which resolves the disharmonic sequence with harmony (the specified repair) will receive a reward (+1). Similarly, for an input containing no disharmony, the faithful candidate is assessed neither a violation nor a reward; a candidate which changes to become disharmonic will receive a violation (–1). But because there is no instance of the marked configuration in the input, no candidate generated from this input can receive the reward for exhibiting the preferred repair.

Wilson (2004, 2006, to appear) shows that TCOT succeeds considerably in accounting for attested harmony patterns, and avoids the pathologies suffered by Agree and Align constraints discussed above. While the formal mechanisms for TCOT and the positive constraint proposed in this dissertation differ, the overall solution is similar. Because both a targeted spreading constraint and a positive spreading constraint prefer a single instance of harmony over no harmony (regardless of the possibility of complete agreement throughout the domain), the Sour Grapes problem of Agree is avoided. And because neither a targeted constraint nor a positive constraint is sensitive to the number of unassimilated segments in a domain, the segment-counting pathological predictions of Align are avoided.

The major point of departure is that the positive constraint proposed here dispenses with the violation-assigning component used by targeted constraints. In TCOT, the negative component of the constraints serves two important functions: allowing targeted constraints to block other processes, and avoiding the infinite goodness problem. Both of these are accomplished in other ways by the positive constraint proposed in this chapter.

Because the reward-assigning component of a targeted constraint like (22) rewards a particular *change*, it can only ever exert a preference for an instance of spreading, not harmony in general. It cannot therefore block other processes that would lead to disharmony; as Wilson points out, the ability for a constraint to both induce changes and block other processes is a crucial insight of OT. However, because the positive constraint Spread(F) assigns rewards to the output of an operation (a linked segment), rather than to an instance of that operation itself (autosegmental association), it retains the ability to block (see 2.3.1 for a discussion) — processes which threaten the autosegmental links created by spreading will produce candidates which perform worse on SPREAD(F).

Additionally, as discussed above, an unrestricted GEN means that a reward-assigning constraint will lead to potentially infinitely improving candidates. The violation-assigning component of the targeted constraint ensures that rewards will only be assessed for *repairs*, meaning that candidates can only improve to the degree that they were marked to begin with. However, the restricted GEN in HS limits the power of a positive spreading constraint, enabling it to avoid the infinite goodness problem.

The result is that the system of preferences exerted by a targeted spreading constraint and the positive SPREAD(F) — and their predicted typologies — are similar. The extent to which the predictions of the two proposals differ is an interesting question, and one which merits further investigation.

2.3 Further Predictions

This section discusses some additional consequences of adopting a positively formulated spreading constraint. While it is not sensitive to the number of non-harmonized segments there are in a word, a positive spreading constraint like SPREAD(F) is sensitive to the number of assimilated segments there are. As such, the prediction is that a high-ranked SPREAD(F) could preferentially protect members of the harmony domain from processes that threaten to reduce the number of those segments; this prediction, which seems to be attested, is discussed in Section 2.3.1.

Section 2.3.2 discusses the generality of the HS solution to the infinite goodness problem. What allows SPREAD(F) to avoid infinite satisfaction is the fact that, once all available segments have been assimilated, adding *new* targets requires multiple steps, the first of which is not harmonically improving. This will not necessarily be true for all conceivable constraints, and it will be assumed that CON may contain both positively formulated and negatively formulated constraints.

2.3.1 Segment Protection Effects

Section 2.2.2 above discusses the advantageous fact that a constraint like SPREAD(F) is not sensitive to the number of unassimilated segments in a word — it therefore does not have the power to block processes like epenthesis which affect the segment count of a word. SPREAD(F) is sensitive, on the other hand, to the number of *assimilated* segments — precisely because it assigns a reward for each dependent. It is therefore predicted to be able to block processes which threaten to remove its dependents.

Consider, for example, a process of vowel reduction. The present discussion abstracts away from the particular conditioning factors for reduction — the ad-hoc constraint REDUCE will stand in for whatever constraint or constraints motivate the process. A ranking of REDUCE $\gg$ IDENT(F) is required for reduction along the spreading feature dimension in general; if SPREAD(F) $\gg$ REDUCE, the predicted result is a language where reduction is blocked just in case the targeted segment is a dependent of a harmony domain. This prediction seems a bit strange at first, but a pattern like this is in fact attested in Kera (Pearce, 2008), where harmony blocks reduction along the F_2 dimension.

Kera displays vowel harmony along both the height and back/round dimensions — here I will abstract away from some of the particular restrictions on its application (in particular, directionality, morphological criteria, and prosodic domains), but see Pearce (2003) for further details. Back/round harmony is right-to-left; in (23a), the stem vowel [ɨ] agrees with the suffix vowel with respect to backness and rounding.

(23) **Vowel Harmony in Kera**

a. Back/round harmony

/ciːr-i/ [ciːri] ‘your (f) head’

/ciːr-u/ [cuːru] ‘his head’

b. Height harmony[10]

/gus-ɛ/ [gusi] ‘to buy’

/sɛːn-u/ [siːnu] ‘his brother’

Kera also exhibits a process of vowel reduction — unstressed vowels shorten and reduce. However, Pearce (2008, 2011) reports that vowels in a back/round harmony domain reduce in duration and along the F1 dimension, but *do not reduce along the F2 dimension*. She compares this with vowels which are not members of harmony domains, which reduce along all dimensions. In other words, harmony appears to be protecting vowels from just the kind of reduction that would threaten their status as dependents of the spreading feature.

I will provide a schematic analysis, to demonstrate how such a result follows from the nature of the SPREAD(F) constraint — this is not, of course, intended to be a full analysis of Kera's complex harmony system. The key components of the analysis are as follows: first, a constraint demanding reduction must dominate IDENT constraints for both backness and height, since in non-harmonized vowels reduction occurs along both dimensions. Second, SPREAD(BACK) must dominate that reduction constraint, protecting a vowel's colour features from reduction. Finally, SPREAD(HIGH) must *not* dominate the reduction constraint, since reduction along the F1 dimension is not protected (but must still dominate faithfulness, for height harmony to occur). This ranking is summarized in (24) below.

(24) SPREAD(BACK) ≫ REDUCE ≫ SPREAD(HIGH) ≫ IDENT(BACK), IDENT(HIGH)

[10] Back/round harmony appears to be parasitic on height, but is counterfed by height harmony (in order to participate in back/round harmony, both vowels must be *underlyingly* of the same height). The present discussion abstracts away from this opacity.

A demonstration of how a ranking like the one in (24) results in a language like Kera, with reduction blocked by membership in a harmony domain, is given in the tableaux in (25) and (26). At the first step in (25), the options of reduction (25.1b) and harmony (25.1c) are in direct opposition — and because Spread(–back) dominates Reduce, the candidate with harmony wins. At the second step, reduction (25.2b) and (faithfulness to) harmony (25.2a) are again in opposition — this time because reduction would mean that the reduced vowel would no longer be [–back], and therefore would no longer earn a reward as a dependent of the feature. Reduction here is blocked for the same reason that it was delayed at the first step — Spread(–back) outranks Reduce. Harmony is preserved, and the derivation converges.

(25) **Step 1**

[+][–] cɨr-i	Spread(–back)	Reduce	Ident(back)
a. [+][–] cɨr-i	W	–1	L
b. [+] cɨr-ə	W	L	–1
c. ☞ [–] cir-i	+1	–1	–1

Step 2: *Convergence*

[–] cir-i	Spread(–back)	Reduce	Ident(back)
a. ☞ [–] cir-i	+1	–1	
b. [–] cir-ə	W	L	–1

Height, on the other hand, is not subject to such protection. At the first step in (26), harmony and reduction are again in competition, for the same reasons as above — however, unlike Spread(–back), Spread(+high) ranks below Reduce. Candidate (26.1b), with reduction, is preferred over (26.1c), with height harmony. At the second step, the same preference is maintained, and the faithful candidate with reduction is selected as the winner. The derivation converges.

(26) **Step 1**

[+][−] gus-ɛ	Reduce	Spread(+high)	Ident(high)
a. [+][−] gus-ɛ	W–1		L
b. ☞ [+] gus-ə			–1
c. [+] gus-i	W–1	L+1	–1

Step 2: *Convergence*

[+] gus-ə	Reduce	Spread(+high)	Ident(high)
a. ☞ [+] gus-ə			–1
b. [+] gus-i	W–1	L+1	–1

To summarize, positive and negative spreading constraints make opposite predictions about which segments are relevant for assessing the well-formedness of harmony. Negative constraints assign violations for disharmonic segments, and so predict that the number of those segments should be influential; positive constraints, on the other hand, assign rewards for harmonized segments, and thus predict that it is the number of *those* segments which should be relevant. The absence of languages instantiating the patterns discussed in Section 2.2.2, and the existence of a language like Kera which exhibits segmental protection effects, suggest that the predictions made by positive constraints are a better fit for the attested typology.

2.3.2 The Generality of Positivity

In Section 2.2.1, it was demonstrated that a positive spreading constraint avoids the infinite goodness problem in HS because of its limited, finite candidate set. However, a finite candidate set is not in and of itself sufficient to prevent infinite goodness. Because HS derivations only converge when the single changes produced by Gen fail to produce harmonic improvement, a derivation must be sufficiently gradual to prevent a positive constraint from inducing an infinite loop of instant satisfaction.

With a constraint like Spread(F), potential satisfaction is finite because it is not possible to both insert a segment and autosegmentally link it at the same time. Consider, however, a context-free markedness constraint like ObsOral, which assigns a reward to every oral obstruent. Whether a constraint of this type is viable as a positive constraint depends crucially on the degree to which epenthesis is gradual.

If a segment and its features can be inserted simultaneously, a constraint like ObsOral can be instantly satisfied by epenthesis. At each successive step in the derivation, yet another oral obstruent can be epenthesized; because ObsOral will continue assigning rewards to non-faithful candidates, the derivation will never converge.

However, McCarthy (2008a) presents evidence from cluster simplication suggesting that deletion of segments and deletion of features constitute separate operations. Maintaining a separation of operations on features and operations on segments in Gen for insertion as well as deletion means that derivations can be sufficiently gradual to prevent constraints like ObsOral from triggering infinitely long derivations.

Consider a derivation like the one in (27), which begins with an input containing a nasal obstruent. At the first step, denasalizing that obstruent earns a reward on ObsOral, at the expense of faithfulness. In order for ObsOral to earn further rewards at the second step, Gen would need to provide a candidate with epenthesis of a fully-specified oral obstruent. However, with separate insertion operations for segments and features, such a candidate is not a possible member of the candidate set. The best that Gen can do is to epenthesize a featureless default segment; this is not sufficient to earn a reward from ObsOral, and gratuitously violates faithfulness. The faithful candidate is the winner, and the derivation converges.

(27) **Step 1**

tap̃	ObsOral	Ident(Nas)	Dep
a. tap̃	W+1	L	
b. ☞ tap	+2	–1	

Step 2: *Convergence*

tap	ObsOral	Ident(Nas)	Dep
a. ☞ tap	+2		
b. tapʔ	+2		W–1

It is not possible to satisfy a constraint like ObsOral through gradual epenthesis; the conditions necessary to earn rewards on even a context-free markedness constraint cannot be created wholesale in one step. At the second step in (27), there is no access to the information that the epenthesized default segment could subsequently become specified as oral and earn a reward on ObsOral. Because each step in an HS derivation must be harmonically improving to prevent immediate convergence, a two-step progression towards further satisfying a constraint is impossible.

Further research is needed to determine the full extent to which all constraints are viably positivie, but it's possible to describe the conditions under which goodness will be infinite — a positive constraint may be infinitely satisfied if the structure it prefers can be created with a single step of a single operation, where that operation is not limited by existing structure. Relevant operations include segment epenthesis, epenthesis of floating features, and prosodic parsing (but only to the extent to which recursive or empty prosodic structure is permitted by Gen).

A constraint like Seg (assign +1 for every segment) would be obviously problematic, as would a constraint like High (assign +1 for every high tone, whether or not it's associated with a tone bearing unit). There appears to be no evidence in favor of the existence of constraints like this, which are in many ways the positive counterparts to constraints like *Seg — see e.g. Gouskova (2003) for arguments against *Seg and constraints like it.

Potentially more difficult to contend with is a constraint like Align(wd,ft)L/R (assign +1 for every prosodic word aligned to the L/R edge of some foot). Constraints like this have been posited to account for some secondary stress systems, but could lead to infinite parsing of recursive prosodic words. One solution is to disallow recursive parsing, but alternative analyses of the handful of stress systems that seem to require such a constraint may be possible.

2.4 Sources of Non-Participation

The Kasem example described in Section 2.2 above represents, in many ways, the platonic ideal of vowel harmony. The vowel inventory of Kasem is given in (28) — all vowels in the inventory are harmonically paired. That is, all vowels have a [+atr] variant and a [−atr] variant. All vowels participate fully in harmony — all vowels in the root are able to initiate a harmonic alternation in the affix, and all affix vowels are able to undergo harmonic alternations.

(28) **Kasem Vowel Inventory**

i				u
	ɪ		ʊ	
e				o
	ɛ		ɔ	
		ʌ		
			a	

The same can be said of tongue root harmony in Baiyinna Orochen, whose vowel inventory is given in (29) — each [+atr] vowel has a harmonic pair, a vowel which differs only in its tongue root specifications. Furthermore, all vowels in Baiyinna

Orochen participate in tongue root harmony, both as a static restriction on vowel combinations within roots and as a productive pattern of alternations (Li, 1996).

(29) **Baiyinna Orochen Vowel Inventory**

i, ii				u, uu
	ɪ, ɪɪ		ʊ, ʊʊ	
				o,oo
			ɔ, ɔɔ	
		ʌ, ʌʌ		
			a, aa	

For example, the possessive affix surfaces as [-ŋi] with [+ATR] roots, but as [-ŋɪ] with [−ATR] roots (30a). Similarly, the dative suffix (30b) undergoes ATR harmony. Baiyinna Orochen also has a process of rounding harmony, which does not involve all vowels in the inventory; in (30c) the derivational suffix undergoes rounding harmony, but regardless of whether the stem and suffix vowels are round or unround, it also undergoes ATR harmony.

(30) **Baiyinna Orochen ATR Harmony** (data from Li 1996)

a. Possessive: [-ŋi]∼[-ŋi]

[−ATR] Roots		[+ATR] Roots	
kɔɔxan-ŋɪ	'child'	bʌjʌ-ŋi	'person'
ɔlɔ-ŋɪ	'fish'	ɲoɲxo-ŋi	'bear'

b. Dative: [-du]∼[-dʊ]

[−ATR] Roots		[+ATR] Roots	
bɪra-dʊ	'river'	urʌ-du	'mountain'
ɔrmɔk-tʊ	'swamp'	owon-du	'pancake'

c. Derivational suffix: [-ksʌ]∼[-ksa]∼[-kso]∼[-ksɔ]

[−ATR] Roots		[+ATR] Roots	
sʊlaxɪ-ksa	'fox hide'	bʌjun-ksʌ	'elk hide'
tɔrɔxɪ-ksɔ	'wild boar hide'	ɲoɲxo-kso	'bear hide'

In both Kasem and Baiyinna, contrasts between [+ATR] and [−ATR] are phonemic for all vowels. Harmony thus involves a categorical alternation between otherwise contrastive phonemes; the segments involved in these harmonic alternations are full participants. This kind of idealized harmony is certainly attested, at least for tongue root harmonies — other languages with full participation of all segments include Degema (Krämer, 2003; Pulleyblank et al., 1995), word-internal harmony in Vata (Kaye, 1982), Diola Fogni (Bakovic, 2000; Krämer, 2003), and the Tungusic languages Even, Literary Nanaj, Ola Lamut, and Solon (Li, 1996). However, full participation is not the norm, and there are many instances in which vowels are prevented from undergoing harmony.

In this section, I discuss some of the causes behind certain segments' failure to participate in harmony processes. Section 2.4.1 discusses the most common, and most widely discussed, reason for non-participation — feature co-occurrence restrictions. Section 2.4.2 discusses how lexical exceptionality (including exceptional behavior of certain lexical strata) can cause non-participation, and finally Section 2.4.3 discusses how faithfulness asymmetries can be responsible for certain segments failing to undergo harmony. This section is primarily concerned with the factors that contribute to segments becoming non-undergoers — for other factors that can prevent harmony (including trigger asymmetries and parasitic restrictions), see Chapter 3.

2.4.1 Feature Co-occurrence Restrictions

In Kasem and Baiyinna Orochen, where all vowels participate in tongue root harmony, each vowel in the language's inventory has a harmonic counterpart. This is, of course, not always the case. In Yoruba, for example, only mid vowels are paired with respect to ATR; high vowels may only be [+ATR], and the low vowel may only be [–ATR] (Archangeli and Pulleyblank, 1994; Bakovic, 2000; Casali, 2008; Pulleyblank, 1996; Orie, 2001, 2003). The inventory is given in (31).

(31) **Yoruba Vowel Inventory**

i				u
e				o
	ɛ		ɔ	
		a		

Yoruba exhibits a process of right-to-left ATR harmony (Archangeli and Pulleyblank, 1994, and others). To the left of a [–ATR] mid or low vowel, only [–ATR] mid vowels (and the [–ATR] low vowel) are permitted (32) — disharmonic words like [oja] or [ɔle] are prohibited.

(32) **Yoruba ATR Harmony: Non-high Vowels**

	[+ATR]		[–ATR]
a. ebe	'heap for yams'	e. ɛsɛ	'foot'
b. ole	'thief'	f. ɔbɛ	'soup'
c. epo	'oil'	g. ɛpa	'groundnut'
d. owo	'money'	h. ɔja	'market'

However, the unpaired low and high vowels are permitted to co-occur with either [+ATR] or [–ATR] mid vowels (33). When the [–ATR] low vowel [a] co-occurs with a [+ATR] vowel, or when a [+ATR] high vowel co-occurs with a [–ATR] vowel, harmony fails to occur.

(33) **Yoruba Disharmony: High and Low Vowels**

	Low Vowels			High Vowels	
a. ate	'hat'	(*ʌte)	c. ilɛ	'land'	(*ɪlɛ)
b. awo	'plate'	(*ʌwo)	d. itɔ	'saliva'	(*ɪtɔ)

The failure of low vowels and high vowels to undergo harmony in Yoruba is correlated with the absence of harmonic counterparts in the inventory. Both can be seen as arising from independent facts about the markedness of feature combinations — antagonistic combinations like [+LOW]/[+ATR] and [+HIGH]/[−ATR] are dispreferred, and a language-wide ban on marked configurations would produce both the absence of those vowels in the inventory and the prevention of harmony where that would be the result.

This correlation between inventory facts and non-participating segments has been noted throughout the literature on vowel harmony (see e.g. Vago 1976; Archangeli and Pulleyblank 1994; Kiparsky and Pajusalu 2003, and many others). This is in part because it is so frequently observed — in the majority of languages with non-participating segments, non-undergoers are precisely those segments which are unpaired in the inventory.

The common approach since Kiparsky (1981), which I follow here, is to derive both inventory facts and non-participation from the general markedness of the absent feature combinations. In Yoruba, for example, one relevant constraint is *(+HI,−ATR), which penalizes the articulatory and acoustically antagonistic combination. If this constraint dominates both SPREAD(−ATR) and IDENT(ATR), the result is a language where high vowels must always be +ATR, even when that means preventing harmony from occurring.

The tableau in (34) demonstrates this. Candidate (34a) is faithful, with a high [+ATR] vowel followed by a mid [−ATR] vowel. It incurs no violations of *(+HI,−ATR) and no violations of IDENT(ATR), but it also earns no rewards on SPREAD(−ATR). Candidate (34b) does earn a reward on SPREAD(−ATR), but at the cost of incurring a violation of higher-ranked *(+HI,−ATR). Harmony is thus prevented from taking place; the faithful candidate is the winner, and the derivation converges.

(34) **Step 1**: *Convergence*

[+][−] i l ɛ	*(+HI,−ATR)	SPREAD(−ATR)	IDENT(ATR)
a. ☞ [+][−] i l ɛ			
b. [+][−] ɪ l ɛ	W–1	L+1	W–1

The effects of *(+HI,–ATR) are observed elsewhere in Yoruba, as well. The high [–ATR] vowels [ɪ] and [ʊ] do not occur in any native words in the language (as reflected in the inventory in 31 above) — furthermore, loanwords containing [ɪ] or [ʊ] are repaired, becoming the [+ATR] counterparts [i] and [u], respectively (Archangeli and Pulleyblank, 1994). For example, the English word 'cook' [kʊk] is borrowed as [kúkù]; the ranking necessary for this is shown in the tableau in (35). The faithful candidate (35a) incurs a violation of *(+HI,–ATR); candidate (35b) repairs that violation, incurring a violation of the lower-ranked IDENT(–ATR), and is selected as the winner.

(35) **Step 1**

kʊk	*(+HI,–ATR)	IDENT(ATR)
a. kʊk	W–1	L
b. ☞ kuk		–1

Step 2: *Convergence*[11]

Any language where a feature co-occurrence restriction like *(+HI,–ATR) prevents harmony will also be a language where the marked feature combination is generally absent. The ranking necessary for harmony is one where SPREAD(F) dominates IDENT(F), and the ranking necessary for feature co-occurrence to prevent harmony is one where *(F,X) dominates SPREAD(F); by transitivity, this means that *(F,X) will dominate IDENT(F), resulting in neutralization. This explains why the two

[11]Convergence with this particular constraint set occurs at Step 2 — in actuality, other tonal and phonotactic adjustments must also be made.

phenomena are so frequently related — non-participation due to feature co-occurrence implies the absence of the marked harmonic pair in the inventory.

However, the absence of a harmonic pair in the contrastive inventory of a language is neither a necessary nor sufficient condition for non-participation. The implication above does not hold when factors other than feature co-occurrence are responsible for the failure of harmony (see Sections 2.4.2 and 2.4.3 for a discussion of those factors). Furthermore, the implication is not bidirectional — the absence of a harmonic pair does not necessarily mean non-participation.

Harmony can result in categorical phonological alternations that produce segments not otherwise present in the language's inventory. For example, in Kinande, only high vowels may be independently [+ATR]. In monosyllabic words, (or words consisting of vowels from only a single height class), the permissible inventory is as in (36) — high vowels may be either [+ATR] or [−ATR], but mid and low vowels may only be [−ATR] (Archangeli and Pulleyblank, 1994). The range of possible vowels in monosyllabic roots is shown in (37) — both [+ATR] and [−ATR] high vowels are permitted (37a,g), but only [−ATR] non-high vowels are possible (37b-f).

(36) **Kinande: Contrastive Inventory**

i			u
	ɪ	ʊ	
	ɛ	ɔ	
		a	

(37) **Root vowels in Kinande**

a. ɛ-ri-liːb-a ‘to cover’

b. ɛ-rɪ-lɪːm-a ‘to cultivate’

c. ɛ-rɪ-hɛːk-a ‘to carry’

d. ɛ-rɪ-kaːr-a ‘to force’

e. ɛ-rɪ-bɔːh-a ‘to tie’

f. ɛ-rɪ-hʊːm-a ‘to beat’

g. ɛ-ri-huːk-a ‘to cook’

Kinande exhibits right-to-left ATR harmony; to the left of a [+ATR] vowel, only [+ATR] vowels are permitted, and both stem and affix vowels alternate. The stems shown above in (37) are shown preceding the high [+ATR] agentive suffix in (38); the [−ATR] stems (37b-g) surface as [+ATR] before [i] (38b-g). Furthermore, harmony iterates into the preceding prefix.[12]

(38) **ATR Harmony in Kinande**

a. ɔ-mu-liːb-i ‘coverer’

b. ɔ-mu-liːm-i ‘farmer (cultivator)’

c. ɔ-mu-heːk-i ‘porter (carrier)’

d. ɔ-mu-kʌːr-i ‘forcer’

e. ɔ-mu-boːh-i ‘tier’

f. ɔ-mu-huːm-i ‘beater’

g. ɔ-mu-huːk-i ‘cook’

The harmonic alternations induced by the agentive suffix in (38) result in surface vowels that are not possible independently — the non-high [+ATR] vowels [e,ʌ,o] in (38c-e) are found *only* as the result of harmony. The surface inventory of Kinande (39) thus departs somewhat substantially from the underlying inventory above.

[12]The leftmost prefix [ɔ-] does not undergo harmony, as it is outside the relevant domain of application — see Archangeli and Pulleyblank (2002) for the role of domains in Kinande harmony.

(39) **Kinande: Surface Inventory**

i					u
	ɪ			ʊ	
e					o
	ɛ			ɔ	
		ʌ			
				a	

The alternations between [e] and [ɛ], [o] and [ɔ], and [ʌ] and [a] are non-contrastive in Kinande — they don't represent a phonemic distinction. However, on the surface, Kinande vowels all participate fully in harmony — all vowels alternate as a result of harmony, and all vowels are able to pass their harmonic value along to a preceding segment. This sort of participation is referred to here as non-contrastive participation, because it involves alternations between segments that are not phonemically contrastive.

There is has been some degree of controversy surrounding the status of the surface [ʌ]∼[a] alternation; some authors have described [a] as a transparent non-participant (Schlindwein, 1987; Steriade, 1987; Mutaka, 1995), while others describe it as participating non-contrastively (Hyman, 1989; Clements, 1990, 1991; Archangeli and Pulleyblank, 1994, 2002). However, Gick et al. (2006) present acoustic and articulatory (ultrasound) data showing that [a] is substantially advanced in [+ATR] contexts. Because the degree of advancement is comparable to that of other harmonic pairs in the language, and because the effect does not diminish over distance from the trigger, they conclude that the alternation is phonological (rather than arising from low-level coarticulation).[13]

[13]See Chapter 4 for further discussion of Gick et al.'s results, and the arguments that the alternations involved are categorical rather than the result of phonetic coarticulation.

Similar non-contrastive participation is found in RTR harmony in Khalkha (Halh) Mongolian (Svantesson et al., 2005). The inventory of the language is given in (40), with the non-phonemic [ɪ] in parentheses. Underlyingly, all vowels except [i] are paired with respect to RTR — [u] with [ʊ], [o] with [ɔ], and [e] with [a].

(40) **Khalkha Mongolian Inventory**

i				u
	(ɪ)		ʊ	
e				o
			ɔ	
		a		

Only vowels with the same RTR specification can co-occur in roots, and suffixes alternate according to the specification of the root vowel. For example, the causative suffix surfaces as [-uɮ] with [−RTR] roots (41a-d), and as [-ʊɮ] with [+RTR] roots (41e-g).

(41) **Khalkha Mongolian RTR Harmony** (data from Svantesson et al. 2005)

a. it-uɮ ‘eat-CAUS’

b. uc-uɮ ‘see-CAUS’

c. xeeɮ-uɮ ‘decorate-CAUS’

d. oɡ-uɮ ‘give-CAUS’

e. ʊʊrš-ʊɮ ‘evaporate-CAUS’

f. jaw-ʊɮ ‘go-CAUS’

g. ɔr-ʊɮ ‘enter-CAUS’

Because [ɪ] cannot occur independently or in initial syllables, and is not considered part of the phonemic inventory of the language, [i] has been described as a non-participant — as such, it has been described as transparent (Kaun, 1995, and others). However, Svantesson et al. (2005) present acoustic data from several speakers of Khalkha showing a systematic difference in the realization of [i] in [+RTR] and [−RTR] harmonic domains. The difference is approximately comparable in magnitude to the difference between [u] and [ʊ], and they conclude that [i] has a [+RTR] allomorph [ɪ] which surfaces in harmony.

The acoustic data from Svantesson et al. suggest that [i] and [ɪ] participate non-contrastively in harmony, rather than [i] behaving as transparent. The accusative suffix [ig]∼[ɪg] surfaces as its [–RTR] variant with [–RTR] root vowels (42a-b), and as its [+RTR] variant with [+RTR] root vowels (42c-d).

(42) **Khalkha Mongolian RTR Harmony**

a.	piir-ig-e	'brush-ACC-REFL'
b.	suuʤ-ig-e	'tail-ACC-REFL'
c.	mʊʊr-ɪg-a	'cat-ACC-REFL'
d.	c^haas-ɪg-a	'paper-ACC-REFL'

The data from Kinande and Khalkha run counter to the claim that harmony is always structure-preserving (Kiparsky, 1985) — it can, in fact, result in segments that are otherwise absent from the inventory. The ranking necessary to produce such a language is one where the feature co-occurrence constraint *(X,F) dominates IDENT(F) (resulting in general neutralization of the contrast), but where SPREAD(F) dominates *(X,F) (forcing participation even where it would create marked segments). This ranking is demonstrated by the tableaux in (43) and (44), using Khalkha as an example.

In (43), the Khalkha grammar is confronted (as it must inevitably be under Richness of the Base) with an input containing a high front [–ATR] vowel [ɪ].[14] The faithful candidate (43a) incurs a violation of the feature co-occurrence constraint *(+HI,–ATR,+FR). Candidate (43b) avoids that violation by advancing the vowel, at the expense of lower-ranked IDENT(ATR), and is selected as the winner. The har-

[14]Unlike in Yoruba, the restriction against high [–ATR] vowels is only enforced for front vowels and not for back vowels in Khalkha.

mony constraint SPREAD(−ATR) exerts no preference one way or the other; because it assigns rewards only to dependent segments, it does not reward (43a), and because the input is monosyllabic, there are no additional vowels to spread to.

(43) **Step 1**

pɪɪr	SPREAD(−ATR)	*(+HI,−ATR,+FR)	IDENT(ATR)
a. pɪɪr		W–1	L
b. ☞ piir			–1

Step 2: *Convergence*

As the tableau in (43) demonstrates, this ranking results in general neutralization of the [ATR] contrast in high front vowels, and the absence of [ɪ] from the inventory in monosyllabic words. The tableau in (44) demonstrates that this ranking will also produce a marked [ɪ] *when it is the result of harmony.* The input is a [−ATR] root, with an underlyingly [i] suffix vowel. Candidate (44a) is faithful; it incurs no reward on SPREAD(−ATR), but neither does it violate *(+HI,−ATR,+FR). Candidate (44b), where the [−ATR] of the root has spread to the suffix vowel, does violate the feature co-occurrence constraint — but it also earns a reward on the higher-ranked SPREAD(−ATR), and hence is selected as the winner. The derivation will converge at the second step — a candidate which changes the suffix vowel back to [i] will lose, because it will relinquish its reward on SPREAD(ATR).

(44) **Step 1**

[−] [+] mʊʊr-ig	SPREAD(−ATR)	*(+HI,−ATR,+FR)	IDENT(ATR)
a. [−] [+] mʊʊr-ig	W	L	L
b. ☞ [−] [+] mʊʊr-ɪg	+1	–1	–1

Step 2: *Convergence*

The factorial typology of the three constraints is given in (45). Where harmony is present, possibilities include languages where all segments undergo harmony, but contrasts with marked segments are neutralized elsewhere, as in Khalkha (45a), languages where all segments undergo harmony and marked segments are not neutralized, as in Kasem (45b), and languages where harmony applies except where it would create a marked segment, and marked segments are neutralized generally, as in Yoruba (45d) — these are all attested.[15]

(45) **Factorial Typology: Harmony and Neutralization**

a. Spread(F) ≫ *(+hi,–atr) ≫ Ident(atr)

Harmony with full participation; neutralization (e.g Khalkha)

b. Spread(F) ≫ Ident(atr) ≫ *(+hi,–atr)

Harmony with full participation; no neutralization (e.g. Kasem)

c. *(+hi,–atr) ≫ Ident(atr) ≫ Spread(F)

No harmony; neutralization (e.g. Hebrew)

d. *(+hi,–atr) ≫ Spread(F)≫ Ident(atr)

Harmony with non-participants; neutralization (e.g. Yoruba)

e. Ident(atr) ≫ Spread(F) ≫ *(+hi,–atr)

Ident(atr)≫ *(+hi,–atr) ≫ Spread(F)

No harmony, no neutralization (e.g. English)

In languages where harmony affects all segments (45ab), there is an additional option for resolving violations of *(X,F) constraints — re-pairing, where a segment is harmonically paired with a counterpart that differs along another feature dimension (Bakovic, 2000). For example, note that in the Khalkha inventory shown in (40)

[15]With the cumulative constraint interaction of Harmonic Grammar — discussed in Section 3.1.1 ahead — there is a possible language in which feature co-occurrence constraints and faithfulness constraints gang up, prohibiting processes from creating marked segments but protecting them underlyingly. This doesn't appear to happen in harmony, but effects of that general type are found elsewhere (see e.g. McCarthy 2003a).

above, [e] has no direct [–ATR] counterpart [ɛ] — however, suffixes with [e] when they are attached to [+ATR] roots do undergo harmonic alternations, but they surface as [a] rather than [ɛ]. This is shown in (46) — the instrumental suffix is [–geːr] with [+ATR] roots (46ab), but [–gaːr] with [–ATR] roots (46de). Likewise, the narrative past suffix is [leː] with [+ATR] roots (46c), but [laː] with [–ATR] roots (46f).

(46) **Khalkha Re-Pairing**

a. guzeː-geːr ‘rumen-INST’

b. deːl-eːr ‘coat-INST’

c. uz-leː ‘see-NARR-PAST’

d. ača-gaːr ‘burden-INST’

e. tʊːlaɪ-gaːr ‘hare-INST’

f. jav-laː ‘go-NARR-PAST’

Re-pairing results when SPREAD(F) ≫ *(X,F) ≫ IDENT(X) — a segment is forced to undergo harmony, despite creating a marked feature combination, but subsequently the feature co-occurrence is resolved by change along the *other* feature dimension. This works somewhat differently in HS than in parallel versions of OT. In parallel, the ranking of SPREAD(F) and *(X,F) would not matter, because the re-paired candidate is available from the start and satisfies both constraints; all that's required is for both SPREAD(F) and *(X,F) to dominate IDENT(X). In HS, though, harmony and re-pairing cannot happen simultaneously — they constitute separate operations. So it's necessary to first harmonize, and then subsequently re-pair — in other words, the marked segment must be temporarily tolerated at an intermediate stage of the derivation, and consequently SPREAD(F) must outrank *(X,F).

This is shown in the derivation in (47). At the first step, the faithful candidate (47.1a) does not incur a violation of the feature co-occurrence constraint *(+FR,–ATR), but earns no rewards on SPREAD(–ATR). Candidate (47.1b) has gratuitously backed the suffix vowel; it earns no rewards on SPREAD(–ATR) and incurs no violations of *(+fr,–atr), but does incur a violation of Ident(fr). The winner, candidate (47.1c), has assimilated the suffix vowel; it violates *(+fr,–atr) and Ident(atr), but earns a reward on the highest-ranked Spread(–atr).

At the second step, the faithful candidate (47.2a) maintains its reward on Spread(–atr), but still incurs a violation of *(+fr,–atr). Candidate (47.2c) resolves the feature co-occurrence violation by changing the suffix vowel to [+atr], but this means it no longer recieves a reward on Spread(–atr). Winning candidate (47.2b) instead resolves the feature co-occurrence violation by backing the suffix vowel — it retains its reward on Spread(–atr) and no longer violates *(+fr,–atr), at the mere cost of a low-ranked Ident(fr) violation.

(47) **Step 1**

[–] [+] jav-leː	Spread(–atr)	*(+fr,–atr)	Ident(atr)	Ident(fr)
a. [–] [+] jav-leː	W	L	L	
b. [–] [+] jav-lʌː	W	L	L	W–1
c. ☞ [–] [+] jav-lɛː	+1	–1	–1	

Step 2

[–] [+] jav-lɛː	Spread(–atr)	*(+fr,–atr)	Ident(atr)	Ident(fr)
a. [–] [+] jav-lɛː	+1	W–1		L
b. ☞ [–] [+] jav-laː	+1			–1
c. [–] [+] jav-leː	W		W–1	L

Step 3: *Convergence*

Re-pairing requires Spread(F) to dominate *(X,F), because the initial step of harmony must be harmonically improving despite the feature co-occurrence violation.

Furthermore, *(X,F) must dominate IDENT(X) to permit change along the alternate feature dimension. This does not, however, necessarily dictate the general repair used when encountering *(X,F) violations in contexts other than harmony — the relative rankings of IDENT(F) and IDENT(X) will determine that.[16]

To summarize, feature co-occurrence restrictions can cause segments to fail to participate in harmony (when their participation would give rise to marked feature combinations); in these languages, a segment's non-participation in harmony is successfully predicted to correlate with the absence of its harmonic counterpart in the inventory. However, absence of a harmonic counterpart is not sufficient grounds for non-participation — languages may force harmony despite markedness conditions otherwise respected (non-contrastive participation) or resolve the marked feature combination by changing along another feature dimension (re-pairing).

2.4.2 Lexical Exceptions

While non-participation in harmony is most frequently attributable to feature co-occurrence restrictions, that is not the only source of non-participation. Even in languages with otherwise-regular harmony processes, there can be exceptional morphemes, whose segments fail to undergo harmony. These segments' failure to participate cannot be attributed simply to markedness — they are featurally indistinguishable from their non-exceptional counterparts, which do undergo harmony. Instead, their exclusion from harmonic alternations is idiosyncratic, and must be lexically specified in some way.

In Turkish, for example, there are regular processes of harmony along the front/back dimension (Lewis, 1967; Clements and Sezer, 1982; Kirchner, 1993; Kaun, 1995). The

[16]In the context of harmony, resolution by change along dimension F is prevented by SPREAD(F). This predicts a language where the repair for *(X,F) violations differs between harmonic contexts and elsewhere — in Turkana, for example, the [+ATR] counterpart to [a] is [o] (*[ʌ]), but English loans with [ʌ] are borrowed with [a] (Noske, 1996).

inventory is as in (48) — high vowels and mid rounded vowels are contrastively paired, with [e] re-pairing to alternate harmonically with [a]. In the discussion that follows, I will abstract away from the re-pairing involved in the [e]∼[a] alternation, for the sake of expositional simplicity, and treat them as though they are direct harmonic counterparts.

(48) **Turkish Vowel Inventory**

i	y	ɨ	u
e	ø		o
		a	

Front and back vowels may not co-occur within a word, and suffixes alternate to take on the harmonic value of the root. For example, the nominative plural surfaces as [-ler] with front vowels (49a–d), and as [-lar] with back vowels (49e–h).

(49) **Turkish Vowel Harmony**

a. ip-ler ‘rope’

b. jyz-ler ‘face’

c. køj-ler ‘village’

d. el-ler ‘hand’

e. pul-lar ‘stamp’

f. sap-lar ‘stalk’

g. son-lar ‘end’

h. kɨz-lar ‘girl’

The pattern in (49) obtains for the majority of affixes in the language. However, the noun-forming suffix [-gen] fails to undergo harmony — it always surfaces with [e], regardless of the harmonic value of the root vowel. In (50cd), we see [-gen] following back vowels, where [-gan] would be expected.

(50) **Turkish: Exceptional Non-Participation**

a. sekiz-**gen**-ler 'octagons'

b. ytʃ-**gen**-ler 'triangles'

c. tʃok-**gen**-ler 'polygons'

d. altɯ-**gen**-ler 'hexagons'

There are a handful of other suffixes in Turkish which also fail to undergo harmony — their non-participation cannot be attributed simply to feature co-occurrence, since the [e] in [-gen] is indistinguishable from the [e] in [-ler], and the featural consequences of harmonizing are the same across morphemes. These affixes are simply exceptions to the general pattern, but this nonetheless gives rise to the presence of non-participating vowels in an otherwise regular system.

The ranking required to produce harmony generally is a ranking of SPREAD(+BACK) ≫ IDENT(BACK) — however, the ranking required to produce the exceptional lexical items is IDENT(BACK) ≫ SPREAD(+BACK). This kind of inconsistency can be resolved by positing lexically-indexed versions of faithfulness constraints (Pater, 2006, 2008b),[17] which may be ranked independently from their general counterparts. So to model the exceptionality of [-gen], the ranking IDENT(BACK)$_{-\text{gen}}$ ≫ SPREAD(+BACK) ≫ IDENT(BACK) is required. This is demonstrated in the tableaux in (51) and (52) below.

In (51), the input contains a root with a back vowel and a regularly-alternating suffix (posited here with an underlying front vowel). The faithful candidate (51a) does not earn any rewards on SPREAD(+BACK), and so loses. The winning candidate (51b) earns a reward on SPREAD(+BACK), at the expense of violating lower-ranked IDENT(BACK) — harmony is the victor in the general case.

(51) **Step 1**

[17] Indexed constraints are not the only way of handling morphological exceptionality; see e.g. Inkelas and Zoll (2007) for a comparison with a theory based on co-phonologies.

[+] [−] son-ler	IDENT(BACK)$_{\text{–gen}}$	SPREAD(+BACK)	IDENT(BACK)
a. [+] [−] son-ler		W	L
b. ☞ [+] [−] son-lar		+1	–1

Step 2: *Convergence*

In (52), the input still contains a back root vowel, as above; however, the suffix is now the exceptionally non-alternating [-gen]. The faithful candidate again earns no rewards on SPREAD(+BACK)). Its competitor (52b) does earn a reward on (SPREAD(+BACK), but this time it comes at the expense of violating the higher-ranked IDENT(BACK)$_{\text{-gen}}$ — the faithful candidate is hence chosen as optimal, and the derivation coverges immediately.

(52) **Step 1**: *Convergence*

[+] [−] tʃok-gen	IDENT(BACK)$_{\text{–gen}}$	SPREAD(+BACK)	IDENT(BACK)
a. ☞ [+] [−] tʃok-gen			
b. [+] [−] tʃok-gan	W–1	L+1	W–1

Lexically-indexed faithfulness provides a satisfactory account of the exceptional non-participation of certain suffixes in Turkish. The case could also be made for modeling these exceptions using indexed markedness constraints — because it is repaired, the [e] in [-gen] must undergo a marked intermediate stage, and ranking an indexed version of this constraint above SPREAD(+BACK) would have the same effect as the indexed faithfulness constraint in the tableaux above. The featural diversity of the handful of exceptional morphemes (see e.g. Finley 2009 and sources cited therein) might indicate the desirability of a single faithfulness constraint indexed to the set of non-participating segments (rather than individual, distinct markedness constraints), as might the fact that the exceptional non-undergoer [-jor] is also protected from the prohibition against non-high round vowels in non-initial syllables (a prohibition ranked sufficiently high in Turkish to cause non-participation in rounding harmony). However, neither of these arguments are insurmountable or overwhelmingly compelling.

A more convincing argument for the role of lexically-indexed *faithfulness* comes from the behavior of loanwords in Finnish, which fail to undergo harmony along the front/back dimension. The Finnish vowel inventory is given in (53) — low vowels and round vowels are all paired harmonically along the front/back dimension, while front unrounded vowels are unpaired.

(53) **Finnish Vowel Inventory**

i	y		u
e	ø		o
æ		a	

In native Finnish roots, vowels must agree in backness; only the unpaired vowels [i] and [e] are exempt from this requirement (54cd).[18] Suffixes alternate to agree with the harmonic value of (participating) vowels in the root — the essive suffix surfaces as [-næ] following front vowels (54a) and as [-na] following back vowels (54bc); the adessive suffix surfaces as [-llæ] following front vowels (54e) and as [-lla] following back vowels (54d).

(54) **Finnish Back Harmony**

[18]The analysis of the non-participation of [i] and [e] is the basic feature co-occurrence analysis presented in Section 2.4.1, with the constraint *(−ROUND,−LOW,+BACK) dominating SPREAD(BACK) and preventing harmony.

a. pøtæ-næ 'table-ESS'

b. pouta-na 'fine weather-ESS'

c. koti-na 'home-ESS'

d. vero-lla 'tame-ADESS'

e. kæde-llæ 'hand-ADESS'

In loanwords, however, front and back vowels may co-occur (55). In [vulgæːri], for example, front [æ:] immediately follows back [u]. The front vowels [y,ø,æ] which fail to alternate in loanwords are not featurally distinct from the same front vowels when they alternate in native words — their only distinction is the lexical stratum to which the word containing them belongs.

(55) **Finnish Disharmonic Loanwords**

a. vulgæːri 'vulgar'

b. tyranni 'tyrant'

c. afæːri 'affair'

d. analyːsi 'analysis'

e. marttyyri 'martyr'

In this case, the indexed constraint preventing harmony must be faithfulness rather than markedness — there is no justified markedness constraint available. In words like [marttyyri], spreading [+BACK] onto [yy] would result in [uu], which is a less-marked segment. Backness and rounding (as in [u]) are mutually acoustically enhancing features, and their co-occurrence is not restricted; frontness and rounding (as in [y]), on the other hand, are acoustically antagonistic. In other words, markedness and harmony are on the same side when it comes to front vowels like [y] and [ø] in Finnish loanwords, so lexically-indexed feature co-occurrence restrictions cannot be responsible for the failure of these segments to participate. Furthermore, even if there was a markedness constraint that preferred [y] to [u], the same transitive ranking we saw in Section 2.4.1 above would mean that we would expect to see the contrast neutralized in loanwords, and it is not.[19]

Faithfulness, then, is the only recourse for preventing harmony in loanwords. The analysis parallels the analysis of Turkish exceptional morphemes above, except that the relevant faithfulness constraint is indexed to the entire stratum of loanwords — a ranking of $\textsc{Ident(back)}_{\text{loans}} \gg \textsc{Spread(+back)} \gg \textsc{Ident(back)}$ produces harmony in the general case (because the spreading constraint dominates general faithfulness) but no harmony in the case of loanwords (because the indexed faithfulness constraint dominates the spreading constraint). The treatment of disharmonic loans in Finnish is particularly relevant to a discussion of the treatment of non-participating segments — see Chapter 5 for further discussion and analysis of this case.

2.4.3 Faithfulness Asymmetries

Lexical exceptionality is not the only way in which faithfulness constraints can lead to particular segments' non-participation in harmony. Perceptually grounded faithfulness asymmetries, as instantiated in the P-Map theory (Steriade, 2001), can also lead to asymmetries in participation; Khalkha rounding harmony represents such a case.

The Khalkha inventory from (40) above is reproduced in (56) below. While the previous discussion of Khalkha centered around its tongue root harmony, there is also a process of rounding harmony, which is of present interest. The inventory contains both high and non-high rounded vowels, and for [+ATR] vowels, backness and rounding covary. In harmony, [a] alternates with [ɔ], and [e] alternates with [o]; [i] and [u] do not alternate.

[19] An alternative to lexically-indexed faithfulness for Finnish would be to appeal to root-specific faithfulness; under such an analysis, harmony in native roots would be a historical artifact, and the disharmonic behavior of loans would indicate that the process was no longer active in stems. Such an analysis would still nonetheless rely crucially on faithfulness, rather than markedness, to interfere with harmony.

(56) **Khalkha Mongolian Inventory**

i				u
	(ɪ)		ʊ	
e				o
			ɔ	
		a		

Harmony is exemplified in (57); the distributive suffix surfaces as [-ʤe] with unround roots (57ab), but as [-ʤo] with round roots (57cd).[20] However, high vowels do not participate in rounding harmony, as shown in (58); the causative suffix surfaces as [-uʤ] regardless of whether the root vowel is round or unround (58a-d). Similarly, the accusative suffix surfaces as [-ig] regardless of the root vowel (58e-h).

(57) **Khalkha Mongolian Rounding Harmony**

a. it-ʤe 'eat-DIST'

b. xeeʤ-ʤe 'decorate-DIST'

c. og-ʤo 'give-DIST'

d. c^{h}oor-ʤo 'decrease-DIST'

(58) **Khalkha Mongolian: High Vowels**

[+ROUND]		[−ROUND]	
a. it-uʤ	'eat-CAUS'	e. piir-ig	'brush-ACC'
b. xeeʤ-uʤ	'decorate-CAUS'	f. teeʤ-ig	'gown-ACC'
c. og-uʤ	'give-CAUS'	g. xoʤ-ig	'foot-ACC'
d. c^{h}oor-uʤ	'decrease-CAUS'	h. poor-ig	'kidney-ACC'

What is interesting about the non-participation of high vowels in Khalkha rounding harmony is that they may be either round or unround; that is, they do not

[20]Because there is also ATR harmony, the distributive suffix actually undergoes a four-way alternation ([ʤo]∼[ʤɔ]∼[ʤe]∼[ʤa]) agreeing with the root for both ATR and rounding.

participate in harmony, but are permitted to contrast with respect to the harmonizing feature. This is not predicted to occur in languages where harmony is prevented by feature co-occurrence; recall from the discussion in Section 2.4.1 that markedness-induced harmony failure implies, by transitivity of ranking, neutralization.

That explanation is not entirely straightforward for Khaklha, though, because alternation between [i] and [u] would involve re-pairing — the direct harmonic counterpart to [i] would be [y], which is in fact absent from the inventory. In light of this fact, it seems a feature co-occurrence explanation might be possible — a ranking of *(+FR,+RD) ≫ SPREAD(+ROUND) ≫ IDENT(FRONT) will produce a language where rounding harmony takes place generally, but fails when it would create a front rounded vowel; the rounding contrast in front vowels is neutralized generally, and front rounded vowels become unrounded. This corresponds to language (45d) in the typology above.

The problem with this analysis is that [e] alternates harmonically with [o]. What is required for this analysis is a ranking of *(+FR,+RD,+HI) ≫ SPREAD(ROUND) (to prevent harmony from creating high front rounded vowels), a ranking of SPREAD(ROUND) ≫ *(+FR,+RD,–HI) (to permit harmony to create non-high front rounded vowels), and a ranking of *(+FR,+RD,–HI) ≫ IDENT(FRONT) (to subsequently eliminate non-high front rounded vowels by making them back). This is demonstrated in the tableaux in (59) and (60) below.

In (59), the input is a root with a round vowel, followed by a suffix with high, front, unrounded [i]. The faithful candidate (59a) earns no rewards for spreading, but neither does it violate the feature co-occurrence constraints. Its competitor does earn a reward on SPREAD(ROUND), but at the fatal cost of violating high-ranked *(+FR,+RD,+HI), and it loses. Because the faithful candidate is chosen as the input, the derivation converges immediately, and [i] has failed to undergo harmony.

(59) **Step 1**: *Convergence*

[+] [–] poor-ig	*(+FR,+RD, +HI)	SPR(R)	*(+FR,+RD, –HI)	ID(R)	ID(F)
a. ☞ [+] [–] poor-ig					
b. [+] [–] poor-yg	W–1	L+1		W–1	

In (60), on the other hand, the suffix vowel is a *non-high* front unrounded [e]. At the first step of the derivation, the faithful candidate (60.1a) again receives neither a reward for spreading nor a violation of either feature co-occurrence constraint. Its competitor (60.1b) earns a reward for spreading, and unlike in (59) above, this reward comes only at the cost of voilating the lower-ranked *(+FR,+RD,–HI), and harmony is the winner. At the second step of the derivation, the faithful candidate (60.2a) and its competitor (60.2b) are tied on SPREAD(ROUND) — however, candidate (60.2b) has changed the suffix vowel to a back [o], eliminating the violation of the feature co-occurrence constraint at the mere cost of violating IDENT(FRONT); it is chosen as the winner, and [e] has successfully been re-paired to [o].

(60) **Step 1**

[+] [–] c^hoor-ʄe	*(+FR,+RD, +HI)	SPR(R)	*(+FR,+RD, –HI)	ID(R)	ID(F)
a. [+] [–] c^hoor-ʄe		W	L	L	
b. ☞ [+] [–] c^hoor-ʄø		+1	–1	–1	

Step 2

[+] [–] c^hoor-ʄø	*(+FR,+RD, +HI)	SPR(R)	*(+FR,+RD, –HI)	ID(R)	ID(F)
a. [+] [–] c^hoor-ʄø		+1	W–1		L
b. ☞ [+] [–] c^hoor-ʄo		+1			–1

Step 3: *Convergence*

The problem with this analysis is that it requires the constraint *(+FR,+RD,+HI) to dominate *(+FR,+RD,–HI) — in otherwords, it requires a ranking under which [y] is considered to be more marked than [ø] (because it violates a higher-ranked markedness constraint). This, however, should not be the case — rounding on non-high vowels is more marked than rounding on high vowels, so [ø] should in fact be the more marked vowel.

The fact that high vowels fail to undergo harmony while mid vowels (which are equally indirectly paired) participate can be explained, without subverting the markedness hierarchy, if faithfulness asymmetries of the type captured by the P-Map theory are taken into consideration.

The central claim of the P-Map theory is that the degree to which a change is unfaithful is related to the perceptual distance involved in that change — a highly percpetible change (in particular, neutralization of a segment with salient perceptual cues to a contrast) is more unfaithful than a less perceptible change (neutralization of a segment with poor cues to a contrast). Steriade (2001) discusses this claim in the context of coda/onset asymmetries in cluster reduction; onsets are more salient carriers of cues to place features than codas, because C-to-V transitions contain more acoustic information about (most) consonantal features than V-to-C transitions. Therefore, deleting a segment from onset position represents a more perceptible change, and is consequently more costly in terms of faithfulness. Formally, this is represented with a hierarchy of faithfulness constraints (either in a fixed ranking or stringency relationships), with more salient cues protected by higher-ranked constraints.

The same reasoning can be applied to inherent rather than positional salience. For example, Linker (1982) shows, for a variety of languages, that high rounded vowels are in fact more round than non-high rounded vowels. Multi-dimensional scaling of listeners' perception of round and unround vowels in Terbeek (1977) shows that this articulatory difference is reflected perceptually — high rounded vowels are perceived as more round than non-high rounded vowels. Therefore, changing the rounding specification of a high vowel reflects greater unfaithfulness (violates a higher-ranked faithfulness constraint) than changing the rounding specification of a non-high vowel.

In Khalkha, then, the non-participation of high vowels is due to the fact that they are protected by higher-ranked faithfulness constraints than their non-high counterparts. The pattern is illustrated by the tableaux in (61) and (62) below.

In (61), the input and candidates are the same as in (59) above. The faithful candidate (61a) incurs no violations, but also earns no reward on spreading. Its competitor (61b) assimilates the suffix vowel, earning a reward on Spread(round) — but because the suffix vowel is high, this incurs a violation of the high ranked faithfulness constraint Id(rd)$_{[+\text{HI}]}$ (in addition to a violation of the lower-ranked, general co-occurrence constraint penalizing front rounded vowels). The faithful candidate wins, and the derivation converges immediately.

(61) **Step 1**: *Convergence*

[+] [−] poor-ig	Id(rd)$_{[+\text{HI}]}$	Spr(rd)	*(+fr,+rd)	Id(rd)	Id(fr)
a. ☞ [+] [−] poor-ig					
b. [+] [−] poor-yg	W−1	L+1	W−1	W−1	

In (62), the input and candidates are the same as in (59) above. At the initial step, the faithful candidate (62.1a) again earns no rewards for spreading, but violates no feature co-occurrence constraints. Its competitor, candidate (62.1b), earns a reward for spreading, and this comes at the cost of violating lower ranked feature co-occurrence constraints and the general version of Ident(round) — it does not violate the higher-ranked Id(rd)$_{[+\text{HI}]}$, because the suffix vowel is non-high. The candidate with harmony is the winner.

At the second step in the derivation, both the faithful candidate and its competitor tie with respect to Spread(back). The faithful candidate (62.2a) retains the marked [ø] suffix vowel, incurring a violation of *(+FR,+RD); candidate (62.2b), on the other hand, changes that suffix vowel to its back counterpart [o], avoiding violating the feature co-occurrence constraint (at the cost of low-ranked IDENT(FRONT)) and is selected as the winner. The derivation will converge at Step 3.

(62) **Step 1**

[+] [−] c^hoor-ʒe	ID(RD)$_{[+HI]}$	SPR(RD)	*(+FR,+RD)	ID(RD)	ID(FR)
a. [+] [−] c^hoor-ʒe		W	L	L	
b. ☞ [+] [−] c^hoor-ʒø		+1	–1	–1	

Step 2

[+] [−] c^hoor-ʒø	ID(RD)$_{[+HI]}$	SPR(RD)	*(+FR,+RD)	ID(RD)	ID(FR)
a. [+] [−] c^hoor-ʒø		+1	W–1		L
b. ☞ [+] [−] c^hoor-ʒo		+1			–1

Step 3: *Convergence*

The faithfulness-based analysis and preceding markedness-based analysis are mechanically the same — some constraint that disprefers harmony on high vowels outranks SPREAD(ROUND), while the constraint that disprefers harmony on mid vowels ranks below it. However, a markedness-driven account requires, in order to account for the non-participation of high vowels but the participation of non-high vowels, a ranking that does not reflect the actual relative markedness of the segments involved. The faithfulness-based analysis, on the other hand, is true to a more general, phonetically grounded system of asymmetries in the degree to which a change is unfaithful.

It is interesting to note that markedness asymmetries and faithfulness asymmetries often pattern in opposite directions, and seem to give rise to opposite effects. Because the rounding contrast is more salient among high vowels than non-high vowels, a change in rounding on high vowels is more unfaithful, and can asymmetrically block rounding harmony (leading to the non-participation of high vowels). On the other hand, because non-high vowels are poor hosts for rounding, non-high rounded vowels are more marked than high rounded vowels; this can asymmetrically block rounding harmony (leading to non-participation of non-high vowels).

This latter possibility is found in rounding harmony in Turkish, whose contrastive inventory is reproduced in (63) below. The non-initial inventory of Turkish, however, is reduced — in particular, non-high round vowels (in parentheses below) are prohibited in non-initial position (Haiman, 1972).

(63) **Turkish Vowel Inventory**

i	y	ɨ	u
e	(ø)		(o)
		a	

Rounding harmony in Turkish is left-to-right and root-controlled; vowels to the right of a rounded vowel must also be rounded (Clements and Sezer, 1982; Kaun, 1995; Kirchner, 1993, and others). This can be seen in (64), where the possessive suffix surfaces as [-im] following a root with an unrounded vowel (64ab) but as [-ym] following a root with a rounded vowel (64cd). However, non-high vowels fail to undergo rounding harmony; for example, the dative suffix in (65) surfaces consistently as [–e], even when the root vowel is rounded (65cd).

(64) **Turkish Rounding Harmony**

a. ip-im 'my rope'

b. ev-im 'my house'

c. syt-ym 'my milk'

d. čøp-ym 'my garbage'

(65) **Turkish Non-High Vowels**

a. ip-e ‘rope–DAT’

b. ev-e ‘house–DAT’

c. syt-e ‘milk–DAT’ *syt-ø

d. čøp-e ‘garbage–DAT’ *čøp-ø

In non-initial syllables, Turkish is essentially just like Yoruba, and captured by a ranking like that in (45d) — harmony is prevented when it would create a marked segment, and the contrast involving that marked segment is (and must necessarily be) neutralized generally. In Turkish, the relevant feature co-occurrence constraint is *(−HI,+RD) (which Kaun 1995 gives the more memorable name *RoLo); the only addition needed to the constraint set to handle the expanded inventory of intial syllables is a high-ranking positional faithfulness constraint protecting them. Thus, a ranking of IDENT(ROUND)$_{\sigma_1}$ $\gg$ *(−HI,+RD) $\gg$ SPREAD(ROUND) $\gg$ IDENT(ROUND) produces a language like Turkish.

Returning again to the seemingly opposite predictions made by markedness and faithfulness asymmetries in rounding harmony, there is an important difference between Turkish and Khalkha. In Turkish, the markedness of non-high rounded vowels prevents harmony from creating them in non-initial syllables, and there is no rounding contrast among non-initial mid vowels. In Khalkha, on the other hand, faithfulness prevents rounding harmony from affecting high vowels, and there *is* a rounding contrast among high vowels.

This is a reflection of the fact that the opposing faithfulness and markedness asymmetries involved do not make purely mirror-image predictions. While they each make opposite (and attested) predictions about which segments will fail to participate in harmony, they differ with respect to the systems of contrast that result when harmony fails — when faithfulness is responsible for blocking harmony, there must necessarily be a contrast for the harmonic feature among the protected class of segments; when markedness is responsible for blocking harmony, there must necessarily be neutralization with respect to the harmonic feature among the prohibited class of segments.

2.4.4 Summary

In this section, I have discussed the various factors which contribute to the non-participation of particular segments with respect to harmony processes. Section 2.4.1 discussed the role of markedness and feature co-occurrence restrictions, by far the most common source of non-participation and the one most frequently discussed in the literature. While non-participation due to markedness restrictions implies absence of a harmonic counterpart in the inventory, I discussed how the absence of a harmonic counterpart is neither a necessary nor sufficient condition for non-participation. Segments without harmonic pairs in the inventory can participate, either be re-pairing (changing along another feature dimension to find a counterpart) or by engaging in categorical alternations with respect to the harmonic feature which do not represent a contrast in the inventory.

Section 2.4.2 discussed the role of lexical exceptions and morphological strata in conditioning non-participation. Segments which ordinarily participate in harmony may find themselves in exceptional morphemes (or an exceptional class of morphemes corresponding to e.g. a loanword stratum in the lexicon) which idiosyncratically fail to participate. This non-participation can be modeled using lexically indexed versions of faithfulness constraints. Furthermore, Section 2.4.3 discusses another role that faithfulness plays in non-participation — perceptually-based asymmetries in degrees of faithfulness can result in protected segments' failure to undergo harmonic alternations.

2.5 Conclusion

In this chapter, I have proposed a theory of vowel harmony in which a positively-formulated constraint drives autosegmental spreading by assigning rewards to segments which assimilate. In Section 20, I showed that this approach avoids the problematic predictions of negatively-defined spreading constraints like AGREE and ALIGN, because a positive spreading constraint is not sensitive to the number of unassimilated segments in a word.

In versions of OT with parallel evaluation, a positive constraint would be untenable — it would encounter the infinite goodness problem, because increasingly many rewards could be gained by inserting infinitely many instances of the preferred structure. I argued that Harmonic Serialism — which has been shown to have independent typological benefits — provides a solution to the infinite goodness problem. Inserting additional instances of the structure preferred by SPREAD(F) would necessitate a two-step process, both inserting an additional segment and autosegmentally linking it to the spreading feature. In HS, those steps must be performed independently, and each must be harmonically improving; because segment insertion is not *a priori* preferred, and because there is no derivational look-ahead, there is no HS derivation which leads to infinite satisfaction of a constraint like SPREAD(F).

Finally, I discussed the range of factors which can interfere with the satisfaction of SPREAD(F) — that is, factors which can cause a segment to fail to undergo harmony. I argue that, while they account for a large number of cases of non-participation, feature co-occurrence restrictions are neither necessary nor sufficient to prevent harmony. Segments may harmonize despite otherwise-obeyed feature co-occurrence restrictions, either engaging in non-contrastive alternations or re-pairing along another feature dimension. Additionally, faithfulness constraints can prevent some segments from participating in harmony — either via indexation to exceptional morphemes or strata, or via perceptually-motivated asymmetries. But while this section has dealt with the factors that may cause non-participation, it has not addressed the treatment of non-participating segments. The question of whether a non-participant will be treated as transparent (skipped over by harmony) or opaque (blocking further propagation of harmony) is taken up in Chapter 3.

CHAPTER 3

TRANSPARENCY, OPACITY, AND TRIGGER COMPETITION

3.1 Introduction

The preceding chapter discussed the formal mechanism driving harmony, and the factors which can lead to a segment's failure to undergo harmony (and its status as a non-participant); this chapter takes up the issue of the treatment of those non-participating segments. In particular, when a non-participant is flanked by two participating vowels, there are two options available — harmony can either skip the intervening non-participant (transparency), or it can be blocked from propagating further (opacity). For example, the Ifẹ and Ọyọ dialects of Yoruba differ with respect to the treatment of non-participants; they are transparent in Ifẹ, but opaque in Ọyọ.

Recall from the discussion in the previous chapter that Yoruba exhibits a process of right-to-left [−ATR] harmony (Archangeli and Pulleyblank, 1994; Bakovic, 2000; Casali, 2008). The inventory is given again in (66) — mid vowels are directly paired with respect to tongue root features, but high and low vowels are not.[1]

(66) **Yoruba Vowel Inventory**

i					u
e					o
	ɛ			ɔ	
		a			

[1]The present discussion will focus exclusively on the behavior of [+ATR] high vowels; see Section 3.3.4 for further discussion of the behavior of low vowels.

High vowels in Yoruba are consistently [+ATR], but can co-occur with both [+ATR] and [−ATR] mid vowels. In the stems in (67), high [i] is in a position to be targeted by [−ATR] harmony, but is prevented from undergoing assimilation by a high-ranked feature co-occurrence constraint against high [−ATR] vowels, as discussed in Chapter 2.

(67) **Yoruba Disharmony: High Vowels**

a. ilɛ 'land' (*ɪlɛ)

b. itɔ 'saliva' (*ɪtɔ)

The Ifẹ and Ọyọ dialects of Yoruba both exhibit the pattern described above. However, in instances where a [+ATR] high vowel is flanked by two mid vowels, the two dialects diverge — in Ifẹ, those mid vowels must agree with respect to [ATR], while in Ọyọ, the initial vowel must agree with the intervening high [+ATR] vowel (Pulleyblank, 1996; Orie, 2001, 2003). In other words, non-participating high vowels are transparent in Ifẹ, but opaque in Ọyọ. This can be seen in the examples in (68) — in the Ifẹ examples on the left, the initial vowel is [−ATR] (agreeing with the final vowel), while in the Ọyọ examples on the right, the initial vowel is [+ATR], agreeing with the intervening non-participant.

(68) **Transparency and opacity in Yoruba**

	Ifẹ	**Ọyọ**	
a.	ɛúrɛ́	eúrɛ́	'goat'
b.	ɛ̀lùbɔ́	èlùbɔ́	'yam flour'
c.	ɔ̀títɔ́	òtítɔ́	'truth'
d.	ɔdídɛ	odídɛ	'parrot'

Furthermore, when a language has multiple non-participating segments, they do not necessarily receive uniform treatment; in Hungarian, for example, non-participating

[i] is consistently treated as transparent, while lower front vowels [e] and [æ] are variably opaque, and differ from one another with respect to the relative frequency of opacity (Hayes and Londe, 2006; Hayes et al., 2009). Additionally, in Hungarian and a number of other languages, a sequence of two transparent vowels behaves as optionally opaque. Languages with asymmetries of this type are especially interesting when brought to bear on the central question of this dissertation (and the particular formal focus of this chapter), namely: what factors influence the choice between transparency and opacity?

There has been a considerable literature devoted to the behavior of non-participating segments, and in particular to transparency in harmony. This is because the interactions involved in transparent harmony are non-local, that is, they involve segments which are not adjacent to one another — the initial [ɔ] in Ifẹ in (68d) agrees with the *final* vowel of the word, despite the fact that another vowel (with the opposing feature value) intervenes between them. This runs counter to the principle of strict locality, the claim that phonological interactions may only ever occur between segments that are adjacent at some level of representation, and considerable attention has been devoted to reconciling transparent harmony with strict locality.

Various strategies have been deployed in analyzing transparent harmony while maintaining strict locality. One approach is to deny transparency altogether, and to argue that segments which seem transparent do in fact undergo harmony (Gick et al., 2006; Gafos, 1998; Benus, 2005; Benus and Gafos, 2007) — however, articulatory and acoustic evidence suggests that in at least some languages, transparent segments are actually transparent. Further discussion of this point is presented in Chapter 4, and formal arguments against specific proposals which seek to eliminate transparency are given in Chapter 6.

Acknowledging the existence of surface non-locality, one strategy for maintaining strict locality of phonological operations is to treat transparent segments as an instance of derivational opacity (Clements, 1977; Walker, 1998; Bakovic, 2000). At the stage at which harmony occurs, the phonlogical operation of harmony takes place strictly locally, and surface-transparent segments are actually undergoers. Those segments are later altered, via a general neutralization process, and only then do they become disharmonic. This approach works well for instances where non-participation is attributable to feature co-occurrence restrictions that are generally respected in the language — however, in cases where transparent segments are harmonically paired, there is no available neutralization process to ensure that the temporary undergoers revert to their original feature value (see in particular the discussion of transparency in Finnish loanwords in Chapter 5).

An alternate approach to preserving strict locality (which ultimately falls victim to the same problem as the derivational opacity approach) is to posit a level of representation at which the segments which interact in transparent harmony are in fact local. This is the fundmental strategy behind the underspecification theory of transparency (Kiparsky, 1981; Archangeli and Pulleyblank, 1994), which posits that segments may be skipped `iff` they are unspecified with respect to the harmonizing feature. Because a neutral vowel is not associated with an autosegment, the autosegments associated with flanking vowels are in some sense adjacent. Again, this approach works well for languages where feature co-occurrence restrictions are responsible for segments' non-participation, but encounters difficulty when segments which are paired with respect to the harmonizing feature are treated as transparent. In this case, because these segments contrast, they must be specified — but if they are specified, they should be unable to be transparent.

Further discussion of both derivational opacity and underspecification approaches is taken up in Chapter 6. Given the tension between strict locality on the one hand and the existence of transparent segments on the other, it is worth re-evaluating the arguments in favor of strict locality; Chapter 4 addresses the question of whether maintaining strictly local interactions provides a more phonetically grounded theory, and Chapter 6 addresses the formal arguments that have been made in favor of strict locality — ultimately, the conclusion is that there is no *a priori* reason to prefer a strictly local analysis, and that the empirical evidence favors a different interpretation of locality.

Theories which depart from the inviolability of strict locality have captured the pressure to be local in a variety of ways. The most common approach, within OT-based frameworks, is to include a violable constraint which penalizes non-local spreading. This approach is able to analyze transparency with a fairly high degree of success, because it permits explicitly non-local representations, but makes a series of unwanted typological predictions. In particular, this kind of approach incorrectly predicts languages which satisfy the locality constraint in ways other than blocking harmony. See Chapter 6 for further analysis of this and other approaches to transparency which involve a relaxation of the strictness of locality.

In the approach pursued in this dissertation, the choice between transparency and opacity represents a competition between potential harmony triggers — the segment bearing the dominant feature value on the one hand, and the local trigger on the other hand. In this approach, opacity is not merely blocking, but spreading of the opposing feature value. Hayes and Londe (2006); Hayes et al. (2009) provide an analysis of this type for transparency/opacity asymmetries in Hungarian vowel harmony, arguing that the difference in non-participating segments' amenability to being treated as transparent could be explained by the extent to which they were preferred as triggers themselves.

The overall approach is shown schematically in (69) — in a V_1–V_2–V_3 sequence, where V_2 is a non-undergoer (and harmony is left-to-right), transparency results when V_1 is chosen as the trigger (69a), and opacity results when V_2 is chosen instead (69b). Because V_1 and V_2 compete to determine which will spread its value of the feature to

V_3, factors which influence their fitness as triggers can influence whether the outcome is transparency or opacity — transparency is more likely when V_1 is a good trigger, and opacity is more likely when V_2 is a good trigger.

(69) **Trigger Competition**

a. *Transparency*

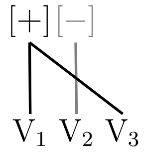

b. *Opacity*

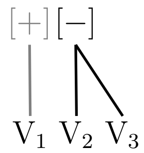

In this chapter, I argue for a version of the competing triggers approach where factors which influence a segment's suitability as a trigger — in particular, locality and quality-specific triggering asymmetries — are encoded as scaling factors on the harmony constraint Spread(F). Rewards for spreading are *increased* from a good trigger, and *diminished* from a poor trigger. The analysis is situated in Serial Harmonic Grammar (Pater, 2008a; Mullin, 2010), a framework which combines the stepwise serial evaluation of Harmonic Serialism (motivated for the present analysis in Chapter 2) with the weighted constraints of Harmonic Grammar (Smolensky and Legendre, 2006).

While the harmony constraint Spread(F) evaluates properties of a surface representation, rewarding segments linked to a feature as dependents, the scaling factors evaluate properties of the spreading operation itself. A segment is a trigger only with respect to a particular instance of autosegmental linking, and does not retain this property. In this way, the present theory successfully accounts for the generalization that the only attested response to conditions on locality or preferred triggers is the failure of harmony — segments never undergo subsequent changes to become better triggers after the fact, and material intervening between target and trigger never deletes or reduces to subsequently minimize the extent of non-locality.

This section is organized as follows. The remainder of this introduction (Section 3.1.1) introduces Harmonic Grammar HG in general and the use of scalar constraints to capture implicational hierarchies in particular. Section 3.2 presents the analysis of non-locality as a diminution of trigger strength, demonstrating how this captures not only transparency and opacity (Section 3.2.1) but also sensitivity to varying degrees of non-locality (Section 3.2.2).

Section 3.3 discusses quality sensitivity in transparency and opacity — that is, languages where some non-participating segments are transparent, and others are opaque. I argue that the differing behavior of non-participants is influenced by independently justified asymmetries in trigger strength (Section 3.3.1), and show how an analysis in which preferred triggers are preferentially opaque accounts for quality sensitivity effects in Hungarian (Section 3.3.2), Menominee (Section 3.3.3) and Wolof (Section 3.3.4). Section 3.4 takes up the issue of similarity sensitivity, and Section 3.5 discusses the specific advantages of the formal architecture proposed here.

3.1.1 Background: (Serial) Harmonic Grammar

Harmonic Grammar (HG) (Smolensky and Legendre, 2006) is a variant of OT in which constraints are assigned numerical weights rather than strict rankings. Each candidate's harmony score is calculated according to the formula given in (70) — the harmony score of some candidate A is the result of multiplying the the number of violations or rewards earned by candidate A on constraint C by the weight of C, and summing across all constraints. The candidate with the highest harmony score is selected as the optimum. See Boersma and Pater (2008); Coetzee and Pater (2008); Jesney (to appear); Pater (2009); Potts et al. (2010) and others for discussion of some of the benefits of modeling constraint interaction with weights rather than strict ranks. Pater (2009), in particular, provides a general overview and introduction to the theory.

(70) $\mathcal{H}(A) = \sum_{i=1}^{n} w(C_i) \times C_i(A)$

The result of combining HG's weighted constraints with HS's restrained GEN and serial evaluation is Serial Harmonic Grammar (SHG). SHG has proven successful at handling a number of challenging patterns, including Berber syllabification (Pater, to appear) and interactions of trigger strength and directionality in Chilcotin (Mullin, 2010).

The difference between OT's ranked constraints and HG's weighted constraints is the potential for cumulative constraint interactions. In OT, domination is strict: if constraint C_1 dominates constraint C_2, the candidate which performs better on C_1 will always win, regardless of how poorly it does on C_2 (or other lower-ranked constraints). In HG, however, it is possible for multiple violations of C_2 to "gang up" on C_1 and result in a different outcome.

Consider the tableaux in (71) and (72). In (71), each candidate violates one of the constraints once — candidate (71a) incurs one violation of Constraint 1, and candidate (71b) incurs one violation of Constraint 2. Because the weight of Constraint 1 is greater than the weight of Constraint 2, candidate (71a)'s violation is more costly, and candidate (71b) is the winner. Under these conditions, weighted constraints behave just like ranked constraints.

(71)

/input/	CONSTRAINT 1 3	CONSTRAINT 2 2	$\mathcal{H}$
a. [output 1]	–1		–3.0
b. ☞ [output 2]		–1	–2.0

In (72), on the other hand, candidate (72b) instead incurs two violations of Constraint 2. The weights are the same as in (71) above — but because candidate (72b)'s harmony score is now the result of multiplying the weight of Constraint 2 by two violations, while candidate (72a)'s harmony score is still the result of multiplying the weight of Constraint 1 by one violation, (72a) comes out ahead — (72b)'s additional violation of Constraint 2 proved too costly. This departs sharply from the result of strictly ranked constraints, where candidate (72b) would still be chosen as optimal.

(72)

/input/	Constraint 1 3	Constraint 2 2	$\mathcal{H}$
a. ☞ [output 1]	–1		–3.0
b. [output 2]		–2	–4.0

Pater (2009); Prince (2004) observe that this kind of cumulative interaction is only possible when there is an asymmetrical tradeoff — that is, when one violation of Constraint 1 can be traded for two or more violations of Constraint 2. In most cases, the exchange will be symmetrical — repairing one markedness violation requires violating one faithfulness constraint, and repairing two markedness violations requires two faithfulness violations. Gang effects like those in (71–72) are therefore restricted to particular contexts.

Scalar constraints — constraints which assign more or fewer marks to different kinds of violation (or satisfaction, in the case of positive constraints) — provide the kind of asymmetrical tradeoff required for cumulative interactions. Pater (to appear) demonstrates that a single scalar constraint in Harmonic Grammar can, by means of cumulativity, replicate the effects of a set of constraints in stringency relationships in OT[2] — see also Flemming (2001). Furthermore, Coetzee (2009) shows that constraint scaling can be used to account for various factors influencing phonological variation.

In this chapter, I propose that the constraint motivating harmony, Spread(F), is in fact a scalar constraint — the rewards for spreading are scaled up or down

[2]Scalar constraints aren't strictly equivalent to constraints in stringency relationships, but in many contexts they make similar predictions. Further research is needed to explore the ways in which they diverge.

according to the locality, quality, and similarity of the target and trigger involved in the operation. When the costs of harmony (faithfuless and markedness) remain constant, the increase or decrease in the reward earned will result in interactions like those in (71) and (72).

3.2 Locality as a Property of Triggers

In the absence of a strict, universal ban on non-local interactions, there must nonetheless be a mechanism available to exert a pressure towards locality. I argue that the preference for locality in harmony is specifically a preference for local triggers — harmony initiated by a segment which is local to the target vowel is preferred over harmony initiated by a non-local trigger. While the constraint driving harmony rewards a particular type of representational structure (a dependent segment associated with an autosegmental feature), a segment's identity as a trigger is operational rather than representational. The definition of the spreading operation from Chapter 2 is reproduced in (73) below, and the definitions of target and trigger are given.

(73) For a feature $\pm F_i$ and a segment S, create an autosegmental association between F_i and S. If S is associated with another instance $\pm F_j$, the link between S and $\pm F_j$ is dissolved. For each instance of an operation creating a new autosegmental link (spreading)...

a. The target is a segment which is not associated with F in the input, but is associated with F in the output (candidate).

b. The trigger is the segment already associated with F in the input which is linearly closest to the target.[3]

[3] See Section 3.4.1 for the desirable consequences of defining the trigger as the linearly closest F-associated segment to the target.

Note that a trigger is distinct from a head — a segment's status as a trigger is a fleeting property defined with respect to a particular instance of the spreading operation, while a segment's status as a head is a durable property of the autosegmental representation. See Section 3.5 for a discussion of how the transience of trigger status prevents conditions on triggers from producing effects other than the failure of harmony.

The triggers and targets of a spreading operation are represented visually in (74). In (74a), +F is associated with S_1 in the input, and an autosegmental linking operation is performed which associates it with S_2 as well — this means that S_2 is the target of the operation (indicated with an underline). Because S_1 is the closest segment already associated with F to the target (in this case it is the only segment associated with F), it is the trigger (indicated with boldface). In (74b), the linking operation creates an association between F and S_3 — S_3 is the target of this operation. Because S_2 is the closest F-associated segment to the target, it is the trigger. Note crucially that the trigger and target identities that S_1 and S_2 took on for the operation in (74a) no longer obtain — they are defined with respect to the operation itself, and as such are not persistent properties of the representation.

(74) **Triggers and Targets**

a. Trigger = S_1, Target = S_2

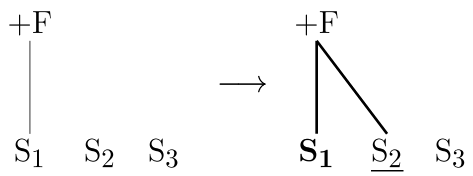

b. Trigger = S_2, Target = S_3

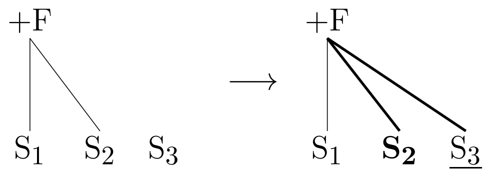

These formal definitions of trigger and target make it possible to make direct reference to the factors that make particular triggers particularly strong (that is, particularly likely to be successful in fulfilling the harmony imperative) — either through inherent properties of the segments themselves or through properties of the relationship between trigger and target. These factors are encoded formally by means of a scaling factor — the reward that SPREAD(F) assigns to a dependent segment is increased or decreased according to the desirability of the trigger (or the trigger/target relationship) involved in that instance of the spreading operation.

The generalization under present discussion is that local triggers are better than non-local triggers. Furthermore, increasing degrees of non-locality are increasingly dispreferred; in Hungarian (Hayes and Londe, 2006), Finnish (Campbell, 1980), and Maltese (Peuch, 1978), harmony is obligatory across one transparent vowel, but either optionally or obligatorily blocked across two (or more) intervening vowels. In other words, increasing the distance between target and trigger decreases the effectivenes of that trigger, decreases the likelihood that harmony will be successful. This is encoded formally as a diminution of the reward earned for spreading, proportional to the distance between target and trigger; the scaling factor is defined in (75). For now, the relevant unit of distance d will be defined as a mora.

(75) **Scaling factor: non-locality**

For a trigger α and a target β, multiply the reward earned for the dependent segment β by a constant k (such that $1 > k > 0$) for each unit of distance d intervening between α and β.

The result of this scaling factor is that an instance of local spreading still receives the usual reward of +1, but an instance of spreading across an intervening vowel will receive a *diminished* reward (and the reward for an instance of spreading across two intervening vowels will be diminished further). In the general case, the precise numerical value of k will not be important;[4] what matters is that it is less than one but greater than zero. In other words, the reward for spreading is diminished by multiplying it by a constant fraction for each increase in distance between trigger and target.

Scaling for non-locality is demonstrated with a toy example in (76); the input contains a [+ATR] vowel followed by two [−ATR] vowels, the constant k has been arbitrarily set at 0.5, and the weights for SPREAD(+ATR) and IDENT(ATR) have been arbitrarily set at 3 and 1 respectively (a set of weights that will result in harmony). At the first step, the faithful candidate (76.1a) earns no rewards on SPREAD(+ATR), but incurs no penalties on IDENT(ATR), so its harmony score is 0. Candidate (76.1b) has spread [+ATR] to the medial vowel; it earns a reward on SPREAD(+ATR), and because this instance of spreading is local, it receives the full reward. It also receives a penalty for violating IDENT(ATR); when its rewards and penalties are multiplied by their respective constraint weights and summed, the result is a harmony score of 2.

Finally, candidate (76.1c) has spread [+ATR] to the final vowel, skipping the medial vowel.[5] It earns a reward on SPREAD(+ATR), but because the spreading operation was non-local (the trigger [i] and the target [o] are separated by mora-bearing [ɛ]), its reward is diminished — it only receives a reward 0.5 (rather than the full reward 1 that its competitor receives). It too violates IDENT(ATR), and when its rewards and violations are multiplied by their constraint weights and summed, the result is a harmony score of 0.5 — candidate (76.1b), with a local trigger, wins.

---

[4]Though see the discussion of Finnish in Chapter 5 — when modeling gradient effects in variable data, the precise extent to which non-locality is dispreferred is reflected in the size of the effect of non-locality on the frequency of transparency.

[5]This analysis assumes that GEN can produce representations with crossed association lines; for discussion of this point, see Chapter 6.

At the second step of the derivation, the candidates each receive a reward for SPREAD(+ATR) — while the scaling factors are evaluated based on operations, the marks themselves are assigned based on autosegmental associations, which remain in place. Each autosegmental link also carries with it some static information about the conditions under which it was formed; any scaling that applied at the step where it was selected as the winner will be retained at subsequent steps, so long as the autosegmental link itself is retained (but because the triggering segment is only a trigger when the operation itself applies, trigger conditions are never re-assessed). This differs somewhat from other HS analyses, where no derivational history is retained, but is an important aspect of the present proposal — see Section 3.5 for further discussion.

The only reward received by the faithful candidate (76.2a) is the reward the autosegmental link it has inherited from the input to this step, and it has a harmony score of 3; its competitor (76.2b) spreads [+ATR] to the final [o] — this time the medial [e] is the trigger, and because there are no intervening vowels, the candidate receives the full reward on SPREAD(+ATR). This reward is combined with the reward for [e]'s status as a dependent (again inherited from the input); candidate (76.2b) also violates IDENT(ATR), and multiplying rewards and violations by their constraint weights and summing results in a harmony score of 5.0. Candidate (76.2b) is the winner, and because there are no vowels remaining to spread to, the derivation will converge at Step 3.

(76) **Step 1**

[+] [−] [−] pi-tɛ-kɔ	SPREAD(+ATR) 3	IDENT(ATR) 1	$\mathcal{H}$
a. [+] [−] [−] pi-tɛ-kɔ			0.0
b. ☞ [+] [−] [−] pi-te-kɔ	+1	–1	2.0
c. [+] [−] [−] pi-tɛ-ko	+0.5 (+1 * 0.5)	–1	0.5

Step 2

[+][−] [−] pi-te-kɔ	SPREAD(+ATR) 3	IDENT(ATR) 1	$\mathcal{H}$
a. [+][−] [−] pi-te-kɔ	+1		3.0
b. ☞ [+][−] [−] pi-te-ko	+2	−1	5.0

Step 3: *Convergence*

It's worth noting that non-local harmony still improves performance on SPREAD(F) (unlike constraints of the AGREE type). However, local harmony will always be preferred — because rewards are diminished for a non-local trigger, any language that permits harmony to occur non-locally will necessarily also permit it locally. This reflects a broader typological generalization: non-local harmony implies local harmony (even in languages where harmony is transparent, transparency represents a subset of cases).

The remainder of this section is organized as follows. Section 3.2.1 discusses how the scaling factor for locality interacts with feature dominance to produce transparency and opacity, and Section 3.2.2 discusses how increasing degrees of non-locality can result in opaque behavior in otherwise-transparent systems.

3.2.1 Transparency and Opacity

In the toy example in (76), both the medial vowel and the final vowel were eligible harmony targets — that is, the initial [i] could induce harmony on the medial [ɛ] as a local trigger, or it could induce harmony on the final [ɔ] as a non-local trigger. Given this choice, the local option is always preferred.[6]

[6]The default preference for locality can be usurped in cases where a distant target is preferred. In metaphony systems, for example, stressed vowels are targeted by height harmony, even when this requires skipping an intervening segment.

However, when the medial vowel is not an eligible target — because, for example, it is prevented from assimilating by a high-ranked feature co-occurrence constraint — a trigger like the initial [i] in (76) no longer has the option of being a local trigger. It can either succeed in triggering harmony non-locally (resulting in transparency) or it can fail to trigger harmony (resulting in opacity).

In order to succeed, the harmony from this non-local trigger must overcome faithfulness, but it must also overcome the possibility of harmony triggered by that non-participating medial vowel, which bears the opposite value for the spreading feature. The pressure towards locality favors harmony triggered by the medial vowel; it receives its full reward, because it is local. However, feature dominance favors the original trigger — the harmony imperative of the dominant feature value outweighs the harmony imperative of the non-dominant value.

While the pressure towards locality is represented as a scaling factor on the harmony constraint, feature dominance is represented by specific versions of the constraint. The reason for this approach is that scalar constraints represent implicational heirarchies that cannot be reversed — local harmony is always preferred to non-local harmony — but feature dominance does not always pattern this way. While for some features, only one vowel is ever dominant (e.g. [+ROUND] or [+NASAL]), there are some features for which either value may be dominant — in particular, [±ATR].[7]

The weighting conditions necessary for harmony that is prevented from creating marked segments are equivalent to the necessary ranking seen in Chapter 2, and are given in (77a). The weighting conditions for feature dominance are given in (77b) — here αF represents the dominant feature value, and $-\alpha$F represents the non-dominant value. Both of the weighting conditions in (77) obtain for either transparency or opacity.

[7]It remains necessary, of course, to explain the cases where only one feature value is ever dominant; the issue is not taken up here.

(77) a. *Weighting conditions for blocked harmony*

$w(*(\mathrm{X},\alpha\mathrm{F})) > w(\textsc{Spread}(\alpha\mathrm{F})) > w(\textsc{Ident}(\mathrm{F}))$

b. *Weighting conditions for feature dominance*

$w(\textsc{Spread}(\alpha\mathrm{F})) > w(\textsc{Spread}(-\alpha\mathrm{F}))$

Broadly speaking, the choice between transparency and opacity represents a choice between locality and dominance. The weighting conditions necessary for transparency are given in (77b) — the weight of dominant SPREAD(αF), when diminished by the scaling factor for non-locality, must still be greater than the sum of the the weights of IDENT(F) and non-dominant SPREAD(–αF). If the weight of either one of these constraints[8] is greater than the diminished weight of SPREAD(αF), harmony will be blocked (78b).

(78) a. *Weighting conditions for transparency*

$w(\textsc{Spread}(\alpha\mathrm{F})) * k > w(\textsc{Ident}(\mathrm{F})) + w(\textsc{Spread}(-\alpha\mathrm{F}))$

b. *Weighting conditions for opacity*

$w(\textsc{Ident}(\mathrm{F})) > w(\textsc{Spread}(\alpha\mathrm{F})) * k$ `or`

$w(\textsc{Spread}(-\alpha\mathrm{F})) > w(\textsc{Spread}(\alpha\mathrm{F})) * k$

The tableau in (79) shows how the weighting conditions in (78a) produce a language like Ifẹ Yoruba, with transparent harmony. The faithful candidate (79.1a) incurs no rewards and no violations, for a harmony score of 0. Candidate (79.1b) has spread [–ATR] to the medial high vowel, and earns a reward on SPREAD(–ATR) (earning the full reward, because spreading is local), but also violates the feature co-occurrence constraint (and faithfulness), resulting in a harmony score of –3. The candidate with transparent harmony (79.1c) earns a reward for SPREAD(–ATR), but

[8]In this case, it is sufficient for *either* IDENT(F) or SPREAD(αF) to be greater than the weight of diminished SPREAD(αF) — if either of them is greater, than the sum of their weights will necessarily be greater, but not vice versa.

because the trigger [ɛ] and the target [ɔ] are separated by a moraic vowel, the reward is multiplied by k (again arbitrarily set at 0.5) and results in a harmony score of 3. Finally, the candidate with opaque harmony (79.1d) earns a reward on SPREAD(+ATR) (and earns the full reward, because spreading is local), and violates IDENT(ATR), resulting in a harmony score of 0.5 — transparent (79.1c) is the winner.[9]

At the second step of the derivation, the input contains the association line between [−ATR] and [ɔ] that was formed at the first step. Candidates (79.2a) and (79.2b) preserve that association line; they each earn a reward on SPREAD(−ATR) — because the operation that created this association line met the criteria for scaling for non-locality, this reward is still scaled, and multiplied by 0.5. The faithful candidate (79.2a) earns no other rewards and incurs no violations, for a harmony score of 3. Candidate (79.2b) has spread to the medial [i], gaining an additional reward of SPREAD(−ATR), but also incurring a violation of the high-weighted co-occurrence constraint (and a faithfulness violation), and ends up with a harmony score of 0. Candidate (79c) has spread [+ATR] from the medial [i], and in the process has eliminated the association link between [−ATR] and the initial erstwhile [ɔ]; it does not receive any rewards from SPREAD(−ATR). It does, however, earn a reward from SPREAD(+ATR) for spreading to the initial [o]; it also violates IDENT(ATR), resulting in a harmony score of 0.5. The faithful (transparent) candidate (79.2a) wins for the same reason that that outcome was selected at the previous step — even a diminished reward on SPREAD(−ATR) is preferred over a full reward on SPREAD(+ATR). Because the faithful candidate is the winnner, the derivation converges.

(79) **Step 1**

[9] In this example, the initial input was already [−ATR], meaning that faithfulness leaves opacity at a disadvantage. However, if the initial vowel had been [+ATR] in the input (as it must sometimes be, under Richness of the Base), the outcome would be the same — the transparent candidate (79.1c) would have a harmony score of 2, and the opaque candidate (79.1d) would have a harmony score of 1.5. In other words, faithfulness here is weighted so low that it does not have influence over the outcome.

[−][+][−] ɔdidɛ	*(+HI,−ATR) 8	SPR(−ATR) 6	SPR(+ATR) 1.5	ID(ATR) 1	$\mathcal{H}$
a. [−][+][−] ɔdidɛ					0.0
b. [−][+][−] ɔdɪdɛ	−1	+1		−1	−3.0
c. ☞ [−][+][−] ɔdidɛ		+0.5 (+1 * 0.5)			3.0
d. [−][+][−] odidɛ			+1	−1	0.5

Step 2: *Convergence*

[−][+][−] ɔdidɛ	*(+HI,−ATR) 8	SPR(−ATR) 6	SPR(+ATR) 1.5	ID(ATR) 1	$\mathcal{H}$
a. ☞ [−][+][−] ɔdidɛ		+0.5 (+1 * 0.5)			3.0
b. [−][+][−] ɔdɪdɛ	−1	+1.5 (+1 * 0.5) + 1		−1	0.0
c. [−][+][−] odidɛ			+1	−1	0.5

The tableau in (80), on the other hand, shows how the weighting conditions in (78b) produce a language with opacity, like Ọyọ Yoruba. The candidates, rewards, and violations are identical to those in the tableau in (79) above — the difference is that the weight of SPREAD(+ATR) is increased. It is still lower than the weight of SPREAD(−ATR), so [−ATR] will still be dominant generally. However, it is now high enough that a diminished reward on SPREAD(−ATR) will lose to a full reward on SPREAD(+ATR). Candidate (80d)'s reward on SPREAD(+ATR) now earns it a harmony score of 4, beating the harmony score of 3 that candidate (80c) receives as a consequence of its diminished reward for spreading [−ATR] non-locally. Under these weighting conditions, opacity is the result.

(80) **Step 1**

[−][+][−] ɔdidɛ	*(+HI,−ATR) 8	SPR(−ATR) 6	SPR(+ATR) 5	ID(ATR) 1	$\mathcal{H}$
a. [−][+][−] ɔdidɛ					0.0
b. [−][+][−] ɔdɪdɛ	−1	+1		−1	−3.0
c. [−][+][−] ɔdidɛ		+0.5 (+1 * 0.5)			3.0
d. ☞ [−][+][−] odidɛ			+1	−1	4.0

Step 2: *Convergence*

The difference between Ifẹ (transparency) and Ọyọ (opacity) is essentially a difference in the releative importance of locality and dominance. In Ifẹ, the weights of SPREAD(−ATR) and SPREAD(+ATR) are relatively far apart, so that even when its reward is diminished due to non-locality, spreading the dominant feature is still preferred. In Ọyọ, on the other hand, the weights of the two harmony constraints, while still ordered the same, are relatively close together — this means that maintaining locality emerges as more important, so that when the reward for spreading the dominant feature is diminished by non-locality, spreading the non-dominant feature locally is preferred.

Both transparency and opacity are robustly attested, despite considerable debate about the status of transparent vowels. Transparency can be found in tongue root harmony in a variety of languages, including Lokaa (Casali, 2008; Akinlabi, 2009), Wolof (Archangeli and Pulleyblank, 1994; Bakovic and Wilson, 2000; Krämer, 2003; Casali, 2008), Orochen (Li, 1996), Classical Manchu (Li, 1996), and others. Transparency with respect to front/back harmony is found in Finnish (Kiparsky, 1981; Ringen, 1988; Ringen and Heinämäki, 1999) and in Hungarian (Kiparsky, 1981; Ringen and Vago, 1995; Benus, 2005); transparency with respect to rounding harmony is found in Khalkha (and other dialects of) Mongolian (Kaun, 1995; Svantesson et al., 2005).

Opacity can be found in many languages and with respect to many different spreading features, including tongue root harmony in Pulaar (Krämer, 2003; Archangeli and Pulleyblank, 1994) and Maasai (Cole, 1987; Bakovic, 2000; Casali, 2008), rounding harmony in Turkish (Kaun, 1995, and others), and height harmony in Kikuria (Cammenga, 1994), among numerous others.

3.2.2 Distance Sensitivity

In the examples in Ifẹ and Ọyọ above, transparency involved skipping only *one* non-participating vowel. However, the scaling factor for non-locality introduced in (75) is sensitive to increasing degrees of non-locality — this reflects a typological generalization that increasing distance between trigger and target decreases the likelihood of transparency.

This can be seen, for example, in Hungarian, which has a process of harmony along the front/back dimension (Hayes and Londe, 2006; Kiparsky, 1981; Kontra and Ringen, 1986; Kornai, 1991; Vago, 1976, and others). The Hungarian vowel inventory is given in (81) — non-low round vowels are harmonically paired, and [ɛ] is paired with [ɔ]. Front unrounded [i] and [e:] are unpaired, and do not participate in harmony — this is attributable to a feature co-occurrence constraint against back unrounded vowels.

(81) **Hungarian Vowel Inventory**

Front		Back
i,iː	y,yː	u,uː
eː	ø,øː	o,oː
ɛ		ɔ,aː

Aside from non-participating [i] and [e], root vowels in Hungarian must agree along the front/back dimension, and suffix vowels alternate to agree with the vowels in their root (82a). Non-participating vowels differ with respect to their treatment, but [i] is consistently treated as transparent (82b). However, when a word contains a sequence of more than one [i], that sequence is optionally opaque (82c).[10]

(82) **Harmony in Hungarian**

a. Suffix vowels agree with back/front stems

hɔd	hɔd-nɔk	'army(-DAT)'
kuːt	kuːt-nɔk	'well(-DAT)'
tœk	tœk-nɛk	'pumpkin(-DAT)'
fyst	fyst-nɛk	'smoke(-DAT)'

b. High front [i,i:] is transparent

gumi	gumi-nɔk	'rubber(-DAT)'
pɔpiːr	pɔpiːr-nɔk	'paper(-DAT)'

c. Multiple neutral vowels can be variably opaque

ɔspirin ɔspirin-nɔk ∼ ɔspirin-nɛk 'aspirin(-DAT)'

The weighting conditions needed to produce the Hungarian data in (82ab) — which pattern like the Ifẹ Yoruba examples analyzed above — are reproduced in (83a–c). In this case, the relevant feature co-occurrence constraint is *(-ROUND,+BACK),[11] and the dominant feature value is [+BACK]. To capture the optional opacity of sequences of multiple neutral vowels, a further weighting condition is needed, given in (83d). While the weight of SPREAD(+BACK), even when its reward is diminished by the scaling factor for locality, must still be greater than the weights of IDENT(F) and SPREAD(–BACK), it must be just low enough that diminishing that reward by

[10]See Hare (1990) for a connectionist analysis of this pattern.

[11]In suffixes, the low(er) front vowel [ɛ] is re-paired to [ɔ]. The present discussion will abstract away from that re-pairing.

repeated applications of the scaling factor (for each unit of distance between trigger and target) results in a decrease sufficient for e.g. SPREAD(–BACK) to triumph.

(83) a. *Weighting conditions for blocked harmony*

$w(*(\mathrm{X},\alpha\mathrm{F})) > w(\textsc{Spread}(\alpha\mathrm{F})) > w(\textsc{Ident}(\mathrm{F}))$

b. *Weighting conditions for feature dominance*

$w(\textsc{Spread}(\alpha\mathrm{F})) > w(\textsc{Spread}(-\alpha\mathrm{F}))$

c. *Weighting conditions for transparency*

$w(\textsc{Spread}(\alpha\mathrm{F})) * k > w(\textsc{Ident}(\mathrm{F})) + w(\textsc{Spread}(-\alpha\mathrm{F}))$

d. *Weighting conditions for distance sensitivity*

$w(\textsc{Ident}(\mathrm{F})) > w(\textsc{Spread}(\alpha\mathrm{F})) * k^n$ or

$w(\textsc{Spread}(-\alpha\mathrm{F})) > w(\textsc{Spread}(\alpha\mathrm{F})) * k^n$

...where n is the distance (in moras) between trigger and target.

Ordinary transparency is demonstrated in (84), using an arbitrary set of weights that meet the criteria in (83) and once again setting k arbitrarily at 0.5. The analysis is essentially identical to the analysis of Ifẹ in (79) above. The faithful candidate (84a) incurs no violations or rewards, and the candidate which has spread the dominant feature (here [+BACK]) locally to the medial vowel (84b) incurs a violation of the highly-weighted feature co-occurrence constraint. The choice, then, is between spreading the non-dominant feature locally from the medial vowel (84c) and spreading the dominant feature non-locally (84d). Because the weight of SPREAD(+BACK) is considerably higher than the weight of SPREAD(–BACK), even a diminished reward for non-local spreading of the dominant feature is preferred over local spreading of the non-dominant feature. Candidate (84d), with transparent harmony, is the winner.[12]

[12]Note that the fact that the suffix is back in the input has no bearing on the outcome — first of all because it is not linked as dependent and therefore has no impact on SPREAD(+BACK), and second because faithfulness is weighted so very low (in fact, its weight is zero).

(84) **Step 1**

[+] [−] [+] gumi-nɔk	*(-RD,+BK) 8	SPR(+BK) 6	SPR(−BK) 2	ID(BK) 0	$\mathcal{H}$
a. [+] [−] [+] gumi-nɔk					0.0
b. [+] [−] [+] gumɨ-nɔk	-1	+1		-1	–2.0
c. [+] [−] [+] gumi-nɛk			+1	-1	2.0
d. ☞ [+] [−] [+] gumi-nɔk		+0.5 (+1 * 0.5)			3.0

Step 2: *Convergence*

Distance sensitivity is demonstrated in the tableau in (85) below; the constraints and weights are the same as in (84) above, where they produced transparency as the favored outcome. In this case, however, instead of one medial non-participant, there are two. As above, the faithful candidate (85a) earns no rewards and incurs no violations, and attempting to spread the dominant feature locally to the medial vowel is ruled out by the high weight of the feature co-occurrence constraint it violates. Once again the choice is between spreading the non-dominant feature locally (85c) and spreading the dominant feature non-locally (85d). However, this time the distance involved in transparency is greater, and therefore the diminution of the reward is more severe. Because there are two transparent vowels (two moras) separating the trigger from the target, the reward is multiplied by k twice — resulting in a reward of only 0.25 (compared to the 0.5 reward earned by the transparent candidate above). Because of this more severe diminution, the harmony score of the transparent candidate is lower than the harmony score of the opaque candidate, which still receives its full reward for spreading (–BACK) locally. Candidate (85d) is the winner, and the result is opacity.

(85) **Step 1**

[+] [-][-] [+] ɔspirin-nɔk	*(-RD,+BK) 8	SPR(+BK) 6	SPR(−BK) 2	ID(BK) 0	$\mathcal{H}$
a. [+] [-][-] [+] ɔspirin-nɔk					0.0
b. [+] [−] [+] ɔspɨrin-nɔk	-1	+1		-1	–2.0
c. ☞ [+] [-][-] [+] ɔspirin-nɛk			+1	-1	2.0
d. [+] [-][-] [+] ɔspirin-nɔk		+0.25 (+1 * 0.5 * 0.5)			1.5

Step 2: *Convergence*

Under the weighting conditions in (83), harmony is predicted to be transparent across a single non-participating vowel, but opaque across two or more non-participating vowels, as it is in Hungarian.[13] In particular, the relationship between the weights of SPREAD(+BACK) and SPREAD(−BACK) is important — they are separated by enough distance that a reward on SPREAD(+BACK) multiplied by k is still worth more than a reward on SPREAD(−BACK), but a reward multiplied by k^2 is not. Furthermore, because the further diminution of the reward for spreading across increasing distances is accomplished via repeated applications of the scaling factor for non-locality, there will be in implicational relationship between e.g. spreading across two neutral vowels and spreading across one — just like there is an implicational relationship between non-local and local harmony.

However, comparing the behavior of [ɔspirin] and [pɔpiːr] indicates the need for a more sophisticated measure of phonological distance. Because scaling is based on the number of moras intervening between trigger and target, the long vowel in [pɔpiːr] predicts that it should be have like [ɔspirin] — in both cases, the inital vowel will be separated from the target in the suffix by two moras (three if coda consonants in

[13]In Hungarian, harmony is *optionally* opaque across a series of non-participants — the present discussion abstracts away from this variation.

Hungarian are moraic). However, the neutral vowel in [pɔpiːr] is reliably treated as transparent.

This is not to say that the distinction between short and long vowels is entirely irrelevant for determining transparency — the nonce word experiment on transparency and opacity in Finnish loanwords presented in Chapter 5 found that short vowels are more likely to be treated as transparent than their long counterparts. Furthermore, coda consontants play a role in blocking tongue root harmony in Lango (Woock and Noonan, 1979; Archangeli and Pulleyblank, 1994; Smolensky, 2006) — harmony is preferred across singletons, and dispreferred across clusters. What is needed, then, is a consistent way of delimiting phonological distance that results in a hierarchy like the one in (86).

(86) .CV. $<$ .CVC. ($<$) .CVV. $<$ CV.CV $< \ldots$

One way to accomplish this is to stipulate such a scale, and to apply the scaling factor for each step along the scale. Alternatively, it would be possible to decompose non-locality into two dimensions along which distance is scaled — for example, a constant k applied for each mora intervening between trigger and target, and an additional constant j applied for each intervening syllable. Regardless of the precise measure of phonological distance used in scale, the multiplicative nature of the scaling factor predicts a quadratic curve, and a fairly steep drop-off — this seems to be a good prediction, as there are relatively few languages with transparent harmony, and fewer still which freely permit harmony across multiple neutral vowels.

3.3 Quality Sensitivity

The choice between transparency and opacity is not just sensitive to the degree of non-locality potentially involved — it may also be sensitive to the features of the non-participating segments. In other words, there are languages in which some non-participants are transparent, while others are opaque.

In Hungarian, for example, the high front vowel [i] is consistently treated as transparent, while non-high [e:] and [ɛ] are optionally opaque (Hayes and Londe, 2006; Hayes et al., 2009) — similarly, in Finnish disharmonic loanwords, high front vowels are more likely to be treated as transparent than low front vowels (Ringen and Heinämäki, 1999). In Menominee, low back [a,a:] is transparent to height harmony, but low front [æ,æ:] is opaque (Bloomfield, 1962). In Eastern Cheremis, rounding harmony affects suffixes but not root vowels, and [ə] in roots is transparent to harmony, but all other vowels are opaque (Odden, 1991). In Khalkha Mongolian, high front [i,ɪ] are transparent to rounding harmony, but [u,ʊ] (despite being round themselves) are opaque (Kaun, 1995; Svantesson et al., 2005). Finally, in Wolof, high [i] and [u] are transparent to [±ATR] harmony, while low [a:] is opaque.

In this section, I argue that the quality sensitivity shown in each of these cases is governed by factors influencing the suitability of each of the non-participating vowels as harmony triggers. In Hungarian, Finnish, Menominee, and Eastern Cheremis, the segments which are (more likely to be) transparent are those which constitute poor triggers, while the segments which are (more likely to be) opaque are better triggers. In Wolof, vowel quality interacts with dominance — the transparent vowels bear the non-dominant feature, while opaque vowels bear the dominant feature. For non-dominant non-participants, dominance and locality are at odds, and this results in transparency as a viable option — but for dominant non-participants, dominance and locality do not conflict, and opacity should always be the result. The issue of transparency and opacity in Khalkha is taken up in Section 3.4, as part of a discussion of the role of similarity sensitivity in harmony.

This section is organized as follows. Section 3.3.1 makes the case for asymmetries in trigger strength, discusses the criteria used to determine the degree to which a segment is a good trigger, and demonstrates how scaling provides an analysis of preferential triggering in general. Section 3.3.2 shows how the interaction of preferential triggering and locality produces the attested asymmetries in Hungarian, Finnish, and Eastern Cheremis, while Section 3.3.3 discusses Menominee. Finally, Section 3.3.4 discusses how dominance and locality interact to produce asymmetries of the type seen in Wolof.

3.3.1 Vowel Quality and Preferential Triggering

In a competing triggers analysis, asymmetries in the amenability of particular segments to transparency is a function of the degree to which those segments constitute good harmony triggers. This approach is based on the fundamental assumption that some segments are better triggers than others — that is, some segments preferentially initiate harmony, while others do not. In this section, I discuss the arguments that such an assumption is independently justified, and show how a scaling factor based on trigger strength is useful in explaining trigger asymmetries generally (beyond cases of transparency and opacity).

Kaun (1995) observes that rounding harmony is preferentially triggered by non-high vowels. In Yakut (Krueger, 1962), for example, there is a process of root-controlled rounding harmony, and both high and non-high vowels are possible triggers. However, non-high vowels trigger harmony in a broader range of contexts — non-high vowels can initiate harmony on both high and non-high targets, while high vowels can only initiate harmony on high targets.

This can be seen in the data in (87) and (88); potential triggers are in boldface, and potential targets are underlined. When the trigger (the final vowel of the stem, in the case of suffix alternations) is high and round, suffixes with high vowels surface as round (87a–c). However, under these same conditions, suffixes with non-high vowels surface as unrounded (87de); harmony has failed to occur. On the other hand, when the trigger is a non-high round vowel, both high suffix vowels (88a–c) and non-high suffix vowels (88de) surface as round.

(87) **Yakut Rounding Harmony: High Triggers**

a. mur**un**-u 'nose–ACC'

b. tuːnn**uː**g-u 'window–ACC'

c. tob**u**g-u 'knee–ACC'

d. tuːnn**uː**k-ter 'window–PL'

e. tob**u**k-ka 'knee–DAT'

(88) **Yakut Rounding Harmony: Non-High Triggers**

a. oɣ**o**-nu 'child–ACC'

b. **o**χ-u 'arrow–ACC'

c. boːr**oː**-nu 'wolf–ACC'

d. oh**o**χ-tor 'stoves–PL'

e. s**o**n-ton 'jacket–ABL'

The reluctance of non-high vowels to undergo harmony can be attributed to the same feature co-occurrence restriction that blocks rounding harmony from applying to non-high vowels in Turkish — *(–HI,+ROUND). But what accounts for the difference between high and non-high triggers? The former are unable to overcome the restriction on marked vowels, but the latter succeed in inducing harmony regardless. There are two sources of explanation in this case — non-high triggers are preferentially capable of inducing harmony on non-high targets because trigger and target agree in height (parasitic harmony; see Section 3.4) or because there is something inherent about non-high vowels that makes them better triggers for rounding harmony. Either way, some direct reference must be made to the quality of the triggering vowel in order to capture the generalization in Yakut.

Some additional evidence for the particular status of non-high vowels as preferred rounding harmony triggers comes from an artificial grammar experiment presented in Finley (2008). In this experiment, native English-speaking subjects were trained on a novel root-controlled rounding harmony pattern, with suffix vowels alternating according to the roundness of the root. In the control condition, subjects received both high-vowel stems and mid-vowel stems in the training data; subjects in one test condition were given only high-vowel stems in the training data, while subjects in the other test condition were given only mid-vowel stems in the training data.

Subjects were then tested on items repeated from the training data, new stems of the same type as the training data (for the subjects who were exposed only to high-vowel stems, these were new high-vowel stems, and vice versa) and new stems of a type they did not see in training (for the subjects who were trained on high-vowel stems only, these were mid-vowel stems, and vice versa). Finley found that subjects did not generalize to stems of the new type in either condition; however, subjects generalized to new stems of the same type in the mid-vowel-only condition, but not in the high-vowel-only condition. In other words, subjects were more successful at learning a pattern of rounding harmony triggered by mid vowels than one triggered by high vowels, lending some support to the claim that mid vowels are preferred as rounding harmony triggers.

Other examples of preferential triggering can be found in Lango (Woock and Noonan, 1979; Archangeli and Pulleyblank, 1994; Smolensky, 2006) and Vata (Kaye, 1982), where high vowels preferentially trigger ATR harmony. In Lango, left-to-right harmony can take place across a consonant cluster when the trigger is a high vowel, but may only take place across singleton consonants otherwise. In Vata, harmony can cross word boundaries under certain circumstances; in these cases, high triggers may induce harmony on both high and non-high targets, but non-high triggers may only induce harmony on non-high targets.

Preferential triggering can also be seen in harmony systems involving both consonants and vowels. In Ennemor, nasal harmony may be triggered by nasalized continuants, but not by nasal stops (Hetzron and Marcos, 1966; Mullin, 2010); in Cairene Arabic, pharyngeal harmony triggered by a flap is blocked by opaque [i,j], but harmony triggered by pharyngealized coronal segments is unblocked (Watson, 2002); similarly, in Chilcotin, regressive pharyngeal harmony is unblocked from preferred coronal triggers, but blocked when induced by dorsal triggers (Krauss, 1975; Cook, 1976; Hansson, 2001; Mullin, 2010).

The evidence above favors the inclusion of some mechanism for exerting preference for particular triggers, but what determines whether a segment will be a good trigger or a poor one? Kaun (1995) proposes, based on the evidence from preferential triggering in rounding harmony systems, that segments which are perceptually impoverished with respect to the harmonizing feature will be preferred as triggers.

Harmony confers a perceptual advantage (Suomi, 1983; Kaun, 1995; Gallagher, 2010) — by realizing a feature across multiple segments, it gives the listener multiple opportunities to correctly identify the intended feature value, thus boosting the perceptual salience of weak cues (see Chapter 4 for further discussion of this point, and experimental evidence in favor of treating harmony as perceptually motivated). Therefore, segments with the weakest perceptual cues for the relevant contrast have the most to gain from harmony — they are in the most dire need of the boost in perceptual salience it provides — and should be preferred as triggers.

This prediction is consistent with the preference for non-high triggers in rounding harmony. Linker (1982) found that, across a variety of languages, non-high round vowels are articulated with a less extreme rounding gesture than their high counterparts. Terbeek (1977) finds that this difference in articulation has perceptual consequences; listeners treat high round vowels as more round than non-high round vowels. In rounding harmony, then, the vowels which are the most poorly cued for roundness are those which have been found to preferentially induce harmony.

The same reasoning applies to backness harmony as well. Higher vowels are farther apart along the front/back dimension than lower vowels, articulatorily and acoustically. The front/back contrast, then, should be less salient in lower vowels, which in turn are predicted to be preferred triggers. While independent evidence for preferential triggering is not found in front/back systems, the behavior of non-high vowels in transparency/opacity asymmetries patterns as would be expected if they are better triggers.

The prediction that perceptually impoverished segments are preferred triggers might appear, at first glance, to be at odds with the status of high [+ATR] vowels as strong triggers in tongue root harmony systems. The features [+HIGH] and [+ATR] are mutually enhancing; high vowels are the best [+ATR] vowels, and are articulatorily and acoustically more advanced than their non-high counterparts (see e.g. Archangeli and Pulleyblank 1994 for a discussion). However, what matters for the present discussion is not the compatibility of a segment's features, but rather the salience of the contrast involved. Because the features [+HIGH] and [–ATR] are articulatorily and acoustically antagonistic, it is likely that they are less perceptually distinguishable than at least their mid counterparts, which would be consistent with their status as preferred triggers. In other words, salience here is not about the inherent strength of a particular cue — the fact that [i] manifests [+ATR] robustly is of little use in distinguishing it from [ɪ], because the articulatory and acoustic difference between the two members of the pair is diminished.

The preference for perceptually impoverished segments as triggers is exerted via the scaling factor in (89). The reward for spreading is increased (multiplied by a constant x which is greater than one) if harmony is induced by a perceptually impoverished trigger. For now, a binary distinction between segments which are perceptually impoverished and segments which are not will suffice, but for each feature, segments can be arranged on a hierarchy according to the salience of the cues for the relevant contrast, and each step on that scale represents a degree of impoverishment.[14]

(89) **Scaling factor: trigger strength**

For a trigger α, a target β, and a feature F, multiply the reward earned for the dependent segment β by a constant x (such that $x > 1$) for each degree i to which α is perceptually impoverished with respect to ±F.

Because the trigger of a given harmony operation is defined with respect to that particular instance of the operation, it is the quality of the derivationally local trigger that is predicted to matter. This is the case, for instance, in Yakut — a word like (87e) [tob**u**k-k<u>a</u>] patterns with other stems with high vowels, failing to induce harmony on the suffix, because the stem-final vowel is the derivationally local trigger when the suffix is targeted for harmony (rather than the initial non-high vowel, which would be successful as a trigger).

Further evidence for derivationally local triggering relationships comes from Hungarian, where a sequence of two neutral vowels is more likely to be treated as transparent when the second neutral vowel is [i] than when it is [e:] (Hayes et al., 2009). See also Jurgec (2011) (and Section 3.4) on icy targets, segments which undergo harmony but which fail to propagate it further because they are unsuitable triggers.[15]

The weighting conditions necessary for preferential triggering are given in (90). The weight of the harmony constraint SPREAD(F) must be such that, when unscaled, it is lower than the feature co-occurrence constraint *(X,F) (or some other constraint

[14]This scale bears some conceptual resemblence to Steriade (2001)'s P-Map — but while the P-Map theory locates the effects of perceptual distance in faithfulness alone, the current proposal claims that markedness may also access this kind of perceptual scale.

[15]Though cf. Walker (2010b), who presents evidence from Baiyinna Orochen and from Jingulu which suggests that the properties of non-local segments might also be relevant.

that could prevent harmony, e.g. faithfulness) — but when multiplied by the scaling factor for preferrential triggering x, the result is greater than the weight of *(X,F).

(90) *Weighting conditions for preferential triggering*

$$w(\text{Spread(F)}) * x > w(\text{*(X,F)}) > w(\text{Spread(F)})$$

Preferential triggering in Yakut is demonstrated in the tableaux in (91–93), with one additional assumption — that high-weighted initial-syllable faithfulness protects the non-high round vowel in the initial syllable, which would otherwise be neutralized by the fact that the weight of the feature co-occurrence constraint, because it outweighs Spread(F) and Spread(F) outweighs faithfulness, must therefore also outweigh faithfulness.

In (91), the input contains a non-high round vowel (in protected initial syllable position) followed by a high unrounded vowel (with a subsequent high unrounded vowel in the suffix). At the first step, each of the relevant candidates incurs a single violation of *(−hi,+round). The faithful candidate (91-1a) incurs no other rewards or violations. Candidate (91-1.b) has spread [+round] to the final high vowel in the suffix; its reward on Spread(+round) is multiplied by the scaling factor for preferential triggering (here arbitrarily set at 2), but is also multiplied by the scaling factor for non-locality. Candidate (91-1.c) has spread [+round] to the medial high vowel; this candidate's reward is multiplied only by the scaling factor for preferential triggering, and not diminished by non-locality. Because the target is high, no additional violations of *(−hi,+round) are incurred — candidate (91c) is the winner, and the initial non-high trigger has succeeded in spreading rounding to a high target.

At the second step, the potential trigger is now the medial high vowel, which is now rounded. Both the faithful candidate (91-2a) and the candidate with rounding spreading to the suffix (91-2b) incur violations of *(−hi,+round) for the initial non-high rounded vowel, and both inherit their reward of +2 on Spread(+round). However, candidate (91-2b) earns an additional reward on Spread(+round) — because the trigger is now the medial high vowel, this reward is not multiplied by the scaling factor for preferential triggering (nor is it multiplied by the scaling factor for non-locality). Regardless, the reward earned is still sufficient to overcome faithfulness; candidate (91-2b) is the winner, and a high trigger has succeeded in spreading rounding to a high target. The derivation will converge on the third step.

(91) **Step 1**

[+] [–] [–] tobɯg-ɯ	*(-HI,+RD) 8	SPREAD(+RD) 6	IDENT(RD) 1	$\mathcal{H}$
a. [+] [–] [–] tobɯg-ɯ	–1			–8.0
b. [+] [–] [–] tobɯg-u	–1	+1 (+1 * 0.5 *2)	–1	–3.0
c. ☞ [+] [–] [–] tobug-ɯ	–1	+2 (+1 * 2)	–1	3.0

Step 2

[+] [–] [–] tobug-ɯ	*(-HI,+RD) 8	SPREAD(+RD) 6	IDENT(RD) 1	$\mathcal{H}$
a. [+] [–] [–] tobug-ɯ	–1	+2 (+1 * 2)		4.0
b. ☞ [+] [–] [–] tobug-u	–1	+3 (+1 * 2) + 1	–1	9.0

Step 3: *Convergence*

Consider, on the other hand, the same stem as in (91), but this time with a non-high suffix.[16] The outcome of Step 1 will be the same — rounding will spread from the initial non-high vowel in the stem to the medial high vowel. The second step, however, will have a different outcome, and is given in (92) below. Both the faithful candidate (92a) and the candidate with harmony on the suffix (92b) again each incur a violation of *(–HI,+ROUND); both candidates still inherit their rewards

[16]The present discussion abstracts away from other unrelated phonological processes that also take place in Yakut.

on Spread(+round), and the candidate with suffix harmony once again incurs an additional (unscaled) reward. However, in this case, the suffix vowel is non-high — spreading [+round] now incurs an additional violation of *(–hi,+round). The penalty for this additional violation is sufficient to prevent harmony from occurring; the faithful candidate is the winner, and the derivation converges. A high trigger has failed to spread rounding to a non-high target.

(92) **Step 2**: *Convergence*

[+] [–] [–] tobuk-ka	*(-hi,+rd) 8	Spread(+rd) 6	Ident(rd) 1	$\mathcal{H}$
a. ☞ [+] [–] [–] tobuk-ka	–1	+2 (+1 * 2)		4.0
b. [+] [–] [–] tobuk-ko	–2	+3 (+1 * 2) + 1	–1	1.0

Finally, in the tableau in (93), the input contains a non-high round stem vowel and a non-high unrounded suffix vowel. Once again, both the faithful candidate (93a) and the candidate with harmony (93b) incur a violation of *(–hi,+round) for the initial non-high round vowel. Candidate (93b) incurs an additional violation, as the suffix is non-high and harmony thus creates an additional non-high rounded vowel. Candidate (93b) also earns a reward on Spread(+round) for harmony — this time, because the trigger is non-high, this reward is multiplied by the scaling factor for preferential triggering. This increased reward is sufficient to overcome the additional violation of the feature co-occurrence constraint, and candidate (93b) is the winner. The non-high trigger has succeeded in spreading rounding to the non-high target (where before the high trigger had failed).

(93) **Step 1**

[+] [−] son-tan	*(-HI,+RD) 8	SPREAD(+RD) 6	IDENT(RD) 1	$\mathcal{H}$
a. [+] [−] son-tan	–1			-8.0
b. ☞ [+] [−] son-tan	–2	+2 (+1 * 2)	–1	–5.0

Step 2: *Convergence*

To summarize, the weighting condition expressed in (90) results in preferential triggering — in the case of Yakut, non-high triggers preferentially induce rounding harmony. While both high and non-high triggers succeed in inducing harmony on high targets — where no violation of *(−HI,+ROUND) is introduced — only preferred non-high triggers are able to overcome the feature co-occurrence restriction to induce harmony when the target is a non-high vowel (and the result of harmony is a marked segment). This is because the weight of the feature co-occurrence constraint is sufficiently high that a single reward on the spreading constraint will not make up for the markedness penalty, but the increased reward earned by a preferred target by virtue of the scaling factor will.

3.3.2 Interactions with Locality: Hungarian

In the preceding chapter, the scaling factor for trigger strength resulted in preferential triggering in general; in this section, I discuss how it can interact with the scaling factor for locality, permitting a preference for particular triggers to influence the choice between transparency and opacity of non-participants.

Section 3.2.2 above discussed the case of Hungarian front/back harmony, where front vowels [i,eː,ɛ] are non-participants. When a neutral [i,iː] intervenes between a back vowel and the suffix vowel, it is reliably transparent — the suffix vowel must surface as back (94).

(94) **High front [i,iː] is transparent**

gumi	gumi-nɔk	‘rubber(-DAT)’	(*gumi-nɛk)
pɔpiːr	pɔpiːr-nɔk	‘paper(-DAT)’	(*pɔpiːr-nɛk)

However, Hayes and Londe (2006); Hayes et al. (2009) identify several zones of variation in Hungarian harmony — contexts where the selection of a front or back suffix becomes variable or unpredictable.

One such zone of variation, discussed above, involves sequences of multiple transparent vowels. Another concerns back–neutral sequences — while [i,i:] are consistently transparent, lower neutral vowels may optionally be opaque. For example, a back–neutral stem like [hotɛl] may surface with either the front version of the dative suffix [-nɛk] or the back version [-nɔk]. Furthermore, a sequence of more than one neutral vowel (even if that neutral vowel is the reliably-transparent [i,i:]) may behave as optionally opaque (95e).

(95) **Hungarian Variable Opacity**

a. ɔrzeːn	ɔrzeːn-nɔk ∼ ɔrzeːn-nɛk	‘arsenic(-DAT)’
b. hotɛl	hotɛl-nɔk ∼ hotɛl-nɛk	‘hotel(-DAT)’
c. honveːd	honveːd-nɔk ∼ honveːd-nɛk	‘Hungarian soldier-DAT’
d. dʒungɛl	dʒungɛl-nɔk ∼ dʒungɛl-nɛk	‘jungle-DAT’
e. ɔspirin	ɔspirin-nɔk ∼ ɔspirin-nɛk	‘aspirin(-DAT)’

In both corpus studies and wug tests, Hayes and Londe (2006); Hayes et al. (2009) found height asymmetries in the probability of transparency in back–neutral sequences. The probability of transparency was highest when the neutral vowel was [i,i:], lower when the neutral vowel was [e:], and lowest of all when the neutral vowel was [ɛ]. This asymmetry is consistent with what would be expected for preferential triggering — lower vowels are less separated along the front/back dimension, and should therefore be more likely to induce harmony by virtue of their perceptual impoverishment. Like the discussion of distance sensitivity above, the present discussion will abstract away from variability, and will treat non-high front vowels as categorically opaque.

The weighting conditions necessary for quality sensitive opacity are given in (96). The conditions in (96a–c) are required for transparent harmony in the ordinary case; it is the additional condition (96d) which is of present interest. While the weight of the constraint favoring spreading the dominant feature is great enough that even a diminished reward for non-local spreading will beat a reward for local spreading of the non-dominant feature, it is not sufficiently high to overcome the increased reward earned by locally spreading the non-dominant feature from a preferred trigger.

(96) a. *Weighting conditions for blocked harmony*

$w(*(\mathrm{X},\alpha\mathrm{F})) > w(\textsc{Spread}(\alpha\mathrm{F})) > w(\textsc{Ident}(\mathrm{F}))$

b. *Weighting conditions for feature dominance*

$w(\textsc{Spread}(\alpha\mathrm{F})) > w(\textsc{Spread}(-\alpha\mathrm{F}))$

c. *Weighting conditions for transparency*

$w(\textsc{Spread}(\alpha\mathrm{F})) * k > w(\textsc{Ident}(\mathrm{F})) + w(\textsc{Spread}(-\alpha\mathrm{F}))$

d. *Weighting conditions for quality sensitivity*

$w(\textsc{Spread}(-\alpha\mathrm{F})) * x > w(\textsc{Spread}(\alpha\mathrm{F})) * k$

Transparency, as demonstrated in (84) above, is reproduced in the tableau in (97) below. The crucial candidates are (97d), which has spread [+BACK] non-locally, and (97c), which has spread [–BACK] locally. Candidate (97d) earns a reward on SPREAD(+BACK), but because spreading is non-local, this reward is multiplied by k and decreased, resulting in a harmony score of 3. Candidate (97c) earns a reward on SPREAD(–BACK) — because this is local spreading, it receives the full reward, but this still results in a harmony score of only 2, and candidate (97d) wins, resulting in transparency.

(97) **Step 1**

[+] [−] [+] gumi-nɔk	*(-RD,+BK) 8	SPR(+BK) 6	SPR(−BK) 2	ID(BK) 0	$\mathcal{H}$
a. [+] [−] [+] gumi-nɔk					0.0
b. [+] [−] [+] gumɨ-nɔk	-1	+1		-1	–2.0
c. [+] [−] [+] gumi-nɛk			+1	-1	2.0
d. ☞ [+] [−] [+] gumi-nɔk		+0.5 (+1 * 0.5)			3.0

Step 2: *Convergence*

The example in the tableau in (98) differs from that in (97) in that its medial non-undergoing vowel is [ɛ], a preferred trigger, rather than [i]. Once again the crucial candidates to compare are (98c), with opaque harmony, and (98d), with transparent harmony. The reward earned by (98d) is the same as above — it earns a reward on SPREAD(+BACK), diminished for non-locality, resulting in a harmony score of 3. Candidate (98c) again receives a reward for SPREAD(–BACK) — however, this time harmony is induced by a preferred trigger [ɛ], so the reward is scaled accordingly, and multiplied by x (here arbitrarily set at 2). This now results in a harmony score of 4 for the opaque candidate (98c), and it is the winner.

(98) **Step 1**

[+] [−] [+] dʒungɛl-nɔk	*(-RD,+BK) 8	SPR(+BK) 6	SPR(−BK) 2	ID(BK) 0	$\mathcal{H}$
a. [+] [−] [+] dʒungɛl-nɔk					0.0
b. [+] [−] [+] dʒungɔl-nɔk	-1	+1		-1	–2.0
c. ☞ [+] [−] [+] dʒungɛl-nɛk			+2 (+1 * 2)	-1	4.0
d. [+] [−] [+] dʒungɛl-nɔk		+0.5 (+1 * 0.5)			3.0

Step 2: *Convergence*

The weighting conditions required to produce transparency of [i] and opacity of [ɛ] are consistent — the specific weights which instantiate them in (97) and (98) are identical — meaning that the two outcomes can co-exist in the same language, as indeed they do in Hungarian. All things being equal, the choice between transparency and opacity — the choice between locality and dominance — will come down in favor of transparency, spreading the dominant feature despite the weakness of non-local triggering. However, when the local trigger is also a preferred trigger for featural reasons (here because it is non-high), that interferes with that choice — the strength of the preferred trigger tips the scales in favor of the opaque outcome.

In the Hungarian examples in (97) and (98), the quality of the non-local trigger is held constant — however, the competing triggers approach predicts that the status of this segment with respect to preferential triggering should also influence the outcome. Categorical alternations of this type are not found in any vowel harmony system,[17] but gradient effects in contexts where transparency and opacity are variable provide an opportunity to test this kind of prediction. Chapter 5 reports the results of a nonce-word study on disharmonic loanwords in Finnish, finding that not only are preferred triggers more likely to be opaque (as seen above), but also that transparency is more likely when the segment spreading the dominant feature value is a preferred trigger. The analysis of the Finnish results, while representing a more detailed set of facts, is at its core fundamentally the same as the analysis of Hungarian above.

An even more extreme example of the interaction of trigger strength and locality comes from Eastern Cheremis, where stems may be disharmonic but suffix vowels alternate according to the backness and rounding of the stem (Odden, 1991). In disharmonic stems, it is the segment closest to the suffix which determines its feature

[17] Though see the consonant–vowel harmony processes mentioned in Section 3.3.1 above.

value — however, [ə] is treated as transparent. If opacity is attributable to a preference for local triggers, then the transparency of [ə] (despite the opacity of all other vowels) is attributable to the fact that it is not specified along the relevant feature dimension,[18] and therefore cannot possibly be a trigger.

3.3.3 Interactions with Locality and Faithfulness: Menominee

Another example of trigger strength asymmetries influencing transparency and locality can be found in Menominee, which has a height harmony process in which long mid vowels are raised to high before high vowels (Bloomfield, 1962, and others). For example, the initial [e:] in [se:pe:w] (99a) becomes [i:] in [si:piah] (99b), and the peninitial [o:] in [otoːtɛːmew] (99c) becomes [u:] in [otu:puakanew] (99d). Low vowels do not participate in harmony, but are not given uniform treatment — the low front vowel [æ:] is opaque, while the low back vowel [a:] is transparent.

(99) **Menominee Raising**

a.	s**eː**peːw	'river'
b.	s**iː**p**i**ah	'river-LOC'
c.	ot**oː**tɛːmew	'he has X as a totem'
d.	ot**uː**hp**u**akanew	'he has X as a pipe'

However, harmony in Menominee has been analyzed both as height harmony and as ATR harmony, and there is no solid consensus in the literature. Bloomfield (1962) describes Menominee as a height harmony system, and this interpretation is adopted by a number of other researchers (Cole and Trigo, 1987, 1988; Cole, 1987; Steriade, 1987). However, Archangeli and Pulleyblank (1994) reanalyze the harmonizing fea-

[18] Odden provides evidence for the interpretation of [ə] as the language's unspecified default vowel — it is, for example, the vowel that results from various reduction processes.

ture as ATR, and this reanalysis is taken up by subsequent authors, including another explicit reanalysis (Milligan, 2000).

Before turning to the issue of transparency and opacity, it is important to clarify which feature is in fact spreading — whether the contrast relevant for determining preferential triggering is height or ATR. At issue primarily is the long vowel inventory, and in particular whether the vowels that Bloomfield transcribes as [eː] and [oː] are true mid vowels (100a) or instead high [−ATR] vowels (100b).

(100) a. **Menominee as Height: Inventory**

iː		uː
eː		oː
	æː	aː

b. **Menominee as ATR: Inventory**

iː			uː
	ɪː	ʊː	
	ʌː		
			aː

Both Archangeli and Pulleyblank (1994) and Milligan (2000) justify their reanalyses primarily on analytical rather than empirical grounds. They argue that blocking of ATR harmony by [ʌː] receives a straightforward explanation, but blocking of height harmony by [æː] is mysterious. This justification crucially depends on the specific set of theoretical assumptions each author works under, and poses no necessary or inherent problem.

Archangeli and Pulleyblank provide some phonetic justification for their feature system by pointing out that Bloomfield describes considerable variation in the realization of *short* mid vowels, including realizations like those in English *pit* and *put* (Bloomfield, 1962, p. 8-9). The variability and quality of these short vowels is of little use in determining the harmonizing feature in Menominee; their status with respect to harmony is unclear, because they are perceptually very weak — they have been described as transparent (Bloomfield, 1962), but it is possible that they are participants. It is the long vowels which clearly participate, so their quality is of more immediate concern than their short counterparts.

With respect to the long vowels, Bloomfield is unequivocal. He describes long vowels as tense, and compares [e:] and [o:] to the vowels in English *pay* and *go* respectively (but notes that in Menominee they are not diphthongized as they are in English). The fact that the short counterparts of these vowels surface as [−ATR] does not mean that the long vowels must themselves be [−ATR].

Archangeli and Pulleyblank also cite the variability of the low front vowel [æ:] as evidence for ATR harmony. Bloomfield describes the range of pronunciation for this vowel as including variants from English *father* to *bad* and from French *brave* to *tête*. This admittedly impressive range of variation is attributed by them to the articulatory conflict between the feature specifications [+ATR] and [+LOW]. However, specifying this low vowel as [+ATR] and [+LOW] does not require positing that [e:] and [o:] are [−ATR], or that the harmonic alternation that [æ:] fails to participate in is one of ATR rather than height.

Milligan (2000) cites discrepancies between Bloomfield's transcriptions and those of other researchers, in particular those in Miner (1975). While there are slight differences concerning some of the vowel qualities, the transcriptions for the long vowels (in particular the vowels [e:] and [o:]) are consistent with Bloomfield's descriptions. Further evidence that [e:] and [o:] are in fact mid comes from a dictionary of Menominee (The Menominee Indian Tribe of Wisconsin, 2005); [e:] is compared to English *bait* and *say*, and [o:] is compared to English *boat* and *so* — the accompanying sound files available on the website (which the skeptical reader is encouraged to listen to) support these comparisons.

Neither Archangeli and Pulleyblank (1994) nor Milligan (2000) provide any compelling empirical evidence to justify the reanalysis of Menominee as ATR harmony, and their analytical arguments are dependent upon the theoretical frameworks in which they are situated. In the absence of any clear reason to abandon previous descriptions, I follow Bloomfield (1962) in treating the harmonizing feature in Menominee as height, rather than ATR.

Returning again to the harmony process, the long vowels [i:] and [u:] trigger height harmony, and the mid vowels [e:] and [o:] participate in harmonic alternations. The status of their short counterparts [e] and [o] is unclear, but short [i] and [u] can initiate harmony, as can be seen in (99) above and in the examples in (101) — segments in bold are participants in harmony; arrowed lines originate from the (initial) harmony trigger, arrows point to undergoers, and transparent vowels are underlined. The examples in (101) below are shown both with short vowels as non-participants, in which case they behave as transparent (left) and with short vowels as participants (right); I will assume henceforth that they are participants.

(101) **Short Vowels in Menominee Harmony**

		participating	
a.	t**oː**ckenæw		'he nudges him'
b.	t**uː**ckeni**i**hæw	t**uː**ck**i**n**i**hæw	'he nudges him in the body/belly'
c.	t**oː**hkopæːhsen		'he lies with buttocks spread'
d.	t**uː**hkop**i**ahnæw	t**uː**hk**u**p**i**ahnæw	'he walks with buttocks spread'
e.	w**eː**nepow		'he dirties his mouth'
f.	new**iː**nep**i**m	n**i**w**iː**n**i**p**i**m	'I dirty his (my?) mouth'

Unlike high and mid vowels, low vowels do not participate in height harmony, and both long and short versions of the low back vowel [a,a:] are transparent to harmony.

In (102), a high [i] in the final syllable causes raising of all preceding mid vowels, regardless of an intervening [a] or [a:].

(102) **Transparent [a,a:] in Menominee**

a. c**eː**paːhkow — ‘he cooks’

b. n**i**c**iː**paːhk**i**m — ‘cook-NOM’

c. s**oː**p**o**maːhkow — ‘he makes sugar’

d. n**i**s**uː**p**u**maːhk**i**m — ‘sugar-maker’

e. seːkaheːˀkow — ‘he puts drops in his eyes’

f. n**i**s**iː**kah**iː**qk**i**m — ‘I put drops in my eyes’

Both short and long versions of the low front vowel [æ,æ:], on the other hand, are opaque to harmony. Examples in (103) are given both as plain transcriptions and in annotated form, with potential participants in bold and connected by a dashed line, and the opaque vowel circled. Eligible mid vowels earlier in the word do not raise, despite the presence of a high vowel trigger. In (103d), height has spread to raise an available mid vowel, but is unable to cross the intervening [æ:] to raise the initial [o:].

(103) **Opaque [æ,æː] in Menominee**

a. soːwaːnæhkiːqsew 'he has his hair blown back by the wind'

s**oː**w<u>aː</u>n æhk**iː**qsew

b. moːnehpæniːw 'he digs potatoes'

m**oː**n**e**hp æn**iː**w

c. peːhtæhkiːˀtaw 'he sticks his head in'

p**eː**ht æhk**iː**ˀtaw

d. poːtawæːtimiːw 'Potawatomi'

p**oː**t<u>a</u>w æː t**i**m**iː**w

f. keːskenæːhcihæw 'he cuts off his fingers'

k**eː**sk**e**n æː hc**i**hæw

In Menominee, the dominant feature value is [+HIGH] — high vowels cause raising of mid vowels, but non-high vowels do not cause lowering of high vowels. The non-participation of low vowels with respect to harmony is attributable to the incompatibility of the features [+HIGH] and [+LOW]. The transparency/opacity asymmetry shown by the two low vowels is in the expected direction, especially given differences in the two segments' realizations. Low back [a,a:] is described by Bloomfield (1962) as being consistently quite low — it is therefore fairly well cued for contrasts along the height dimension. On the other hand, the considerable variability of low front [æ,æ:], and the fact that its range includes realizations that are very nearly mid vowels, mean that it is poorly cued for height. The prediction, then, is that [æ,æ:] should be preferred as a trigger, and should therefore be more likely to be opaque — in this case, categorically opaque, while its well-cued back counterpart is categorically transparent.

The fact that only assimilation to the dominant feature value results in alternations in Menominee — that is, the fact that non-high vowels do not induce lowering of high vowels — would at first glance appear to be at odds with the attribution of [æ,æː]'s opacity to its status as a preferred trigger. It is now that the distinction between IDENT(F) and DEP(LINK) — the two faithfulness constraints potentially violated by the spreading operation — becomes important. IDENT(F) is violated if a segment is associated with αF in the input, but $-\alpha$F in the output — but DEP(LINK) is violated by the creation of any autosegmental association, even if it does not result in a change in the dependent segment's feature value. This difference means that under the weighting conditions in (104) — with IDENT(F) weighted higher than SPREAD(F), which in turn weighted higher than DEP(LINK) — the result will be vacuous linking. In other words, segments will become associated to an autosegmental feature F `iff` they can do so without changing their feature value.

(104) *Weighting conditions for vacuous linking*

$w(\textsc{Ident}(\text{F})) > w(\textsc{Spread}(\text{F})) > w(\textsc{Dep}(\textsc{link}))$

Linking of already-featurally-matching segments has long been a part of the literature on autosegmental theory, and its use here is not novel. Evidence for vacuous linking can be seen elsewhere, for example in the behavior of certain tone sequences with respect to the Obligatory Contour Principle (OCP) — see e.g. Archangeli and Pulleyblank (1994); Myers (1997) and others for examples and discussion.

Vacuous linking in Menominee is demonstrated in the tableau in (105). The input contains two non-high vowels; the faithful candidate (105a) incurs no violations and no rewards. Its competitor has vacuously spread [–HIGH] — that is, an autosegmental link has been formed between the feature [–HIGH] and a segment which was already non-high to begin with, and so undergoes no visible featural change. This earns a reward on Spread(–hi) and violates Dep(link), but does not violate Ident(hi) — candidate (105b) is the winner, and the derivation will converge at the first step.[19]

(105) **Step 1**

[–] [–] seːp-eːw	Spread(+hi) 14	Id(hi) 4	Spread(–hi) 2	Dep(link) 1	$\mathcal{H}$
a. [–] [–] seːp-eːw					0.0
b. ☞ [–] [–] seːp-eːw			+1	–1	1.0

Step 2: *Convergence*

In Menominee, it is the harmony constraint for the non-dominant feature value, Spread(–hi), whose weight matches the criteria for vacuous linking in (104) by being less than the weight of Ident(hi) and greater than the weight of Dep(link). The dominant feature, however, spreads non-vacuously — the weight of its spreading constraint, accordingly, must be greater than that of Ident(hi).

The differing behavior of the dominant and non-dominant features can be seen in the tableau in (106), where the input contains an initial non-high vowel followed by a high vowel and then another subsequent non-high vowel. The faithful candidate (106a) earns no rewards and no violations; candidate (106d), with vacuous linking of the initial non-high vowel across the medial high vowel, ties with the faithful candidate — it earns a scaled down reward for non-local spreading, and incurs a violation of Dep(link).[20] Candidate (106b) has spread [+high] to the initial vowel, and earns a reward on Spread(+hi); this more than compensates for the violations of Ident(hi) and Dep(link), giving (106a) a harmony score of 9.

[19]Evidence for the viability of representations like (105a) can be found in Boyce (1990)'s study of lip-rounding in English and Turkish. Turkish speakers, for whom sequences of multiple round vowels represent a continuous feature domain, showed a plateau through an intervening consonant. English speakers, on the other hand, showed two separate peaks, one for each rounded vowel, suggesting independent rounding gestures (and therefore independent rounding features).

[20]This tie would be broken if the weight of Dep(link) was slightly higher or slightly lower, and the overall outcome would not be affected.

Candidate (106c), on the other hand, has spread [–HIGH] from the final vowel to the medial vowel — this earns it a reward on SPREAD(–HI), but because spreading to the medial vowel changed its feature value, it also incurs a violation of IDENT(HI) in addition to its violation of DEP(LINK), and its harmony score is –3. Candidate (106b), with dominant-feature harmony, is the winner — but even in the absence of that candidate, (106c) would lose; it scores worse than the faithful candidate. Note also that, even if that final non-high vowel had been the preferred trigger [æ], it would still lose — the increased reward would still be less than the cost of faithfulness.

(106) **Step 1**

[–] [+][–] seːp-i ah	SPREAD(+HI) 14	ID(HI) 4	SPREAD(–HI) 2	DEP(LINK) 1	$\mathcal{H}$
a. [–] [+][–] seːp-i ah					0.0
b. ☞ [–] [+][–] siːp-i ah	+1	–1		–1	9.0
c. [–][+][–] seːp-e ah		–1	+1	–1	–3.0
d. [–] [+][–] seːp-i ah			+0.5 (+1 * 0.5)	–1	0.0

Step 2: *Convergence*

When SPREAD(–HI) is subjected to the weighting conditions in (104), the feature will vacuously link to segments which already bear that same feature value. However, they will *only* be able to link vacuously — because the weight of IDENT(HI) is greater, linking that induces any kind of harmonic alternation will be prevented.

In (106), local dominant harmony (106b) competed with local non-dominant harmony (106c), and soundly defeated it. However, when the competition is between non-local dominant harmony and local vacuous linking of the non-dominant feature, the outcome is less certain. Under the conditions in (96a–c) and (104) and with additional weighting conditions in (107), non-dominant harmony will win just in case it is induced from a preferred trigger (multiplied by the scaling factor x) and its competitor is both non-local (multiplied by the scaling factor k) and feature-changing (violating IDENT(F)).

(107) $w(\textsc{Spread}(-\alpha\mathrm{F})) * x > w(\textsc{Spread}(\alpha\mathrm{F})) * k - w(\textsc{Ident}(\mathrm{F})) > w(\textsc{Spread}(-\alpha\mathrm{F}))$

The tableaux in (108) and (109) show how these weighting conditions produce the pattern of transparency and opacity seen in Menominee. In (108), the final high vowel is separated from its eligible harmony targets by a low back [a:]. At the first step, the faithful candidate (108.1a), with no spreading, earns no rewards and incurs no violations; candidate (108.1b), which has targeted the low vowel, incurs a violation of high-weighted feature co-occurrence. The contenders here are (108.1c), where [+HIGH] has spread from the high vowel across the intervening low vowel, and (108.1d), where [–HIGH] has vacuously spread from the low vowel.

Candidate (108.1c) earns a reward on SPREAD(+HI), but this reward is diminished because the trigger and target are non-local. Because the target's feature value has changed, it also incurs a violation of IDENT(HI). Candidate (108.1d) on the other hand receives no IDENT(HI) violation, because the target was already [–HIGH], and it receives a full reward for spreading because trigger and target are local. Despite this inequity, the weight of SPREAD(+HI) is high enough so that candidate (108.1c)'s harmony score is still greater than that of its competitor, and it is selected as the winner. At the second step of the derivation, [+HIGH] easily succeeds in spreading (locally this time, because the trigger is now the peninitial vowel) to the remaining target, and the derivation will converge at the third step.

(108) **Step 1**

[−][−] [−] [+] necеːpaːhkim	*(+HI,+LO) 15	SPR(+HI) 14	ID(HI) 4	SPR(−HI) 2	DEP(LINK) 1	$\mathcal{H}$
a. [−][−] [−] [+] neceːpaːhkim						0.0
b. [−][−] [−] [+] neceːpaːhkim	–1	+1	–1		–1	-6.0
c. ☞ [−][−][−] [+] neciːpaːhkim		+0.5 (+1 * 0.5)	–1		–1	2.0
d. [−][−] [−] [+] neceːpaːhkim				+1	–1	1.0

Step 2

[−][−][−] [+] neciːpaːhkim	*(+HI,+LO) 15	SPR(+HI) 14	ID(HI) 4	SPR(−HI) 2	DEP(LINK) 1	$\mathcal{H}$
a. [−][−][−] [+] neciːpaːhkim		+0.5 (+1 * 0.5)				7.0
b. [−][−][−] [+] neciːpaːhkim		+0.5 (+1 * 0.5)		+0.5 (+1 * 0.5)	–1	7.0
c. ☞ [−][−][−] [+] niciːpaːhkim		+1.5 (+1 * 0.5) + 1	–1		–1	16.0

Step 3: *Convergence*

In (109), the final high vowel is again separated from its potential targets by a low vowel — but this time, that low vowel is low front [æ], which is a better trigger than the intervening [a] was above. At the first step, candidates (109.1a) and (109.1b) fare the same as in (108) above, and the candidates to consider are (109.1c), with non-local harmony of the dominant feature, and (109.1d), with local vacuous linking of the non-dominant feature. As above, candidate (109.1c) earns a reward on SPREAD(+HI), but that reward is diminished because it is non-local; (109.1c) also incurs a violation of Ident(hi), which (109.1d) does not.

Candidate (109.1d) earns a reward on Spread(−hi), but this time harmony (however vacuous) is initiated by a preferred trigger, so its reward is increased by the scaling factor x. This increased reward is enough to tip the scales in its favor this time — candidate (109.1d) is the winner, and harmony from the dominant trigger is blocked by vacuous linking from the non-dominant trigger. At the second step, local vacuous linking (109.2b) is preferred over both the faithful candidate (109.2a) and even more distant harmony from the dominant feature (109.2c). The derivation will converge on the third step.

(109) **Step 1**

[–] [–] [–] [+] mo:nehpæni:w	*(+hi,+lo) 15	Spr(+hi) 14	Id(hi) 4	Spr(−hi) 2	Dep(link) 1	$\mathcal{H}$
a. [–] [–] [–] [+] mo:nehpæni:w						0.0
b. [–] [–] [–] [+] mo:nehpæni:w	–1	+1	–1		–1	–6.0
c. [–] [–] [–] [+] mo:nihpæni:w		+0.5 (+1 * 0.5)	–1		–1	2.0
d. ☞ [–] [–] [–] [+] mo:nehpæni:w				+2 (+1 * 2)	–1	3.0

Step 2

[–] [–] [–] [+] mo:nehpæni:w	*(+HI,+LO) 15	SPR(+HI) 14	ID(HI) 4	SPR(–HI) 2	DEP(LINK) 1	$\mathcal{H}$
a. [–] [–] [–] [+] mo:nehpæni:w				+2 (+1 * 2)		4.0
b. ☞ [–] [–] [–] [+] mo:nehpæni:w				+3 (+1*2)+1	–1	5.0
c. [–] [–] [–] [+] mu:nehpæni:w		+0.25 (+1 * 0.5 * 0.5)	–1	+2 (+1 * 2)	–1	2.5

Step 3: *Convergence*

To summarize, asymmetries in the strength of [a] and [æ] as harmony triggers can explain their differing behavior with respect to height harmony — [a] is transparent because its perceptual cues for height are comparatively salient, and therefore it is a poor trigger, while [æ] is opaque because it is poorly cued with respect to height, and so it is a preferred trigger. Furthermore, these triggering asymmetries can emerge even in systems where only the dominant feature is able to induce harmonic alternations — the non-dominant feature can vacuously link to segments already bearing its same feature value, and when the non-dominant trigger is particularly strong, this outcome can be preferred over non-local and faithfulness-costly harmony from the dominant trigger.

3.3.4 Interactions with Dominance: Wolof

Wolof has a [–ATR]-dominant system of harmony; high [+ATR] vowels [i,u] are transparent, while the long low [–ATR] vowel is opaque. The language's inventory is given in (110). What is particularly interesting about Wolof is that some of the non-participating segments ([i,u]) bear the non-dominant value of the spreading feature, while the other non-participant ([a:]) bears the dominant value.

(110) **Wolof Vowel Inventory**

i, i:				u,u:
e,e:				o
	ɛ,ɛː	ʌ	ɔ,ɔː	
		a,a:		

Vowels within a word must agree with respect to ATR, and suffixes alternate to take on the specification of the root vowels — the suffixes in (111), for example, surface as [–ATR] with [–ATR] root vowels, and as [+ATR] with [+ATR] root vowels.

(111) **Wolof ATR Harmony**

a. reːr-e ‘be lost in’

b. doːr-e ‘hit with’

c. dɛm-ɛ ‘go with’

d. xɔːl-ɛ ‘look with’

e. reːr-oːn ‘wanted’

f. rɛːr-ɔːn ‘had dinner’

High vowels, which have no [–ATR] counterparts, do not undergo harmony. When they intervene between two participating vowels, they are transparent — the final vowels in (112) surface as [–ATR], despite the [+ATR] [i] or [u] immediately preceding. Their non-participation, like the non-participation of [i] and [u] in ATR harmony in Yoruba, can be attributed to a constraint against high [–ATR] vowels.

(112) **Wolof Transparent [i,u]**

a. tɛkkilɛːn ‘untie!’

b. mɔytulɛːn ‘avoid!’ (cf. [gʌstuleːn], ‘do research!’)

c. sɔppiwulɛːn ‘you have not changed’

d. xɔlliwɔːn ‘peeled’

e. tɛːruwɔːn ‘welcomed’

f. yɛbbijina ‘he went to unload’

Long and short low vowels differ with respect to their involvement in harmony. The short low vowel alternates, surfacing as [a] in [−ATR] contexts, but [ʌ] in [+ATR] contexts (113), and is a full participant in the harmony system.

(113) **Wolof Short [a/ʌ] Participation**

a. sofoːr-ʌm 'his/her driver'

b. tɔːl-am 'his/her field'

c. doːr-ʌnte 'to hit each other'

d. xɔːl-ante 'to look at each other'

e. wett-ʌli 'to give company to'

f. fɛcc-ali 'to fill completely'

The long low vowel [a:], however, does not undergo harmony, and is always [−ATR]. Unlike the [+ATR] non-undergoers [i] and [u], however, [a:] behaves as opaque. Vowels to the right of [a:] in (114) are invariably [−ATR], despite initial [+ATR] vowels. (There are also a handful of exceptional morphemes with short [a] which do not alternate; these pattern with [a:], and are opaque).

(114) **Wolof Opaque [a:]**

a. yobbuwaːlɛ 'to carry away also'

b. genaːlɛ 'to go out also'

c. doːraːtɛ 'to hit usually'

d. jeːmʌntuwaːlɛːti 'to try also without conviction once more'

The present discussion will abstract away from the divergent behavior of short and long low vowels — why [a] alternates with [ʌ], but [a:] does not alternate with [ʌ:]. The constraint *[ʌ:] will stand in for the constraint or constraints actually responsible.[21] Of present interest is the differing treatment of high non-participants [i,u] (which are transparent) and the low non-participant [a:] (which is opaque).

[21]One possible explanation: a gang effect betwe:n *(−HI,+ATR) and a constraint like *LONGV.

The difference is attributable to the fact that the transparent segments bear the non-dominant feature, while the specified segments bear the dominant feature. Transparency of [i,u] is handled in exactly the same way as Ifẹ Yoruba, with a set of weights that meet the criteria in (115) — when the constraint preferring the spreading of the dominant feature (in this case, [−ATR]) is greater than the weight of the constraint preferring the spreading of the non-dominant feature, by a wide enough margin that even a diminished reward for spreading the dominant feature non-locally beats a full reward for spreading the non-dominant feature locally.

(115) a. *Weighting conditions for blocked harmony*

$w(*(\mathrm{X},\alpha\mathrm{F})) > w(\textsc{Spread}(\alpha\mathrm{F})) > w(\textsc{Ident}(\mathrm{F}))$

b. *Weighting conditions for feature dominance*

$w(\textsc{Spread}(\alpha\mathrm{F})) > w(\textsc{Spread}(-\alpha\mathrm{F}))$

c. *Weighting conditions for transparency*

$w(\textsc{Spread}(\alpha\mathrm{F})) * k > w(\textsc{Ident}(\mathrm{F})) + w(\textsc{Spread}(-\alpha\mathrm{F}))$

This is demonstrated for Wolof in the tableau in (116). The faithful candidate (116a) earns no rewards and incurs no violations; the candidate which spreads the dominant feature locally (116b) incurs a violation of the high-weighted feature co-occurrence constraint. The choice is between the transparent candidate (116c) and the opaque candidate (116d) — the former earns a diminished reward on SPREAD(−ATR), and the latter earns a full reward on SPREAD(+ATR).[22] Candidate (116c)'s diminished reward still results in a higher harmony score, and transparency is the result.

(116) **Step 1**

[22]The present discussion abstracts away from strength differences between high, mid, and low triggers in ATR harmony.

[−][+] [+] tɛːruwoːn	*(+HI,−ATR) 16	SPR(−ATR) 12	SPR(+ATR) 2	ID(ATR) 1	$\mathcal{H}$
a. [−][+] [+] tɛːruwoːn					0.0
b. [−][+] [+] tɛːrʊwoːn	–1	+1		–1	–5.0
c. ☞ [−][+] [+] tɛːruwɔːn		+0.5 (+1 * 0.5)		–1	5.0
d. [−][+] [+] tɛːruwoːn			+1		2.0

Step 2: *Convergence*

The tableau in (117) shows what happens when low [aː] is the intervening non-participant. Candidates (117a) and (117b) are the same as above; the faithful candidate earns no rewards and incurs no violations, and the constraint with local spreading of the initial vowel incurs a violation of a high-weighted markedness constraint (the relevant markedness constraint has been substituted in for this case). The choice, again, is between the transparent candidate (117c) and the opaque candidate (117d). However, this time it is the transparent candidate which has spread the non-dominant feature, and the opaque candidate which has spread the dominant feature. The opaque candidate gets a full reward on higher-weighted SPREAD(−ATR), while the transparent candidate gets a diminished reward on the already-low-weighted SPREAD(+ATR). Candidate (117d) wins easily, and the outcome is opacity.

(117) **Step 1**

[+] [−] [+] do:ra:te	*[ʌ:] 16	Spr(−atr) 12	Spr(+atr) 2	Id(atr) 1	ℋ
a. [+] [−] [+] do:ra:te					0.0
b. [+] [−] [+] do:rʌʌte	−1	+1		−1	−5.0
c. [+] [−] [+] do:ra:te			+0.5 (+1 * 0.5)		1.0
d. ☞ [+] [−] [+] do:ra:tɛ		+1		−1	11.0

Step 2: *Convergence*

The reason that [a:] is opaque in Wolof, while the other non-participants are transparent, is that in its case there is no longer a competition between conflicting interests of dominance and locality — both forces are in agreement. The opaque candidate in this instance has the benefit both of spreading locally and of spreading the dominant feature, and will win regardless of how the conflict between locality and dominance is resolved elsewhere.

3.4 Similarity Sensitivity

This section discusses two somewhat different ways in which harmony is similarity sensitive — that is, influenced or restricted by the extent to which the segments involved are similar to one another. In the first case (discussed in Section 3.4.1), the similarity that matters is that between a non-undergoer and its potential trigger; in several Mongolian languages, it is high *rounded* vowels which are opaque to rounding harmony. In the second case (discussed in Section 3.4.2), non-participation itself results from dissimilarity between trigger and target; in cases of parasitic harmony, spreading of one feature only takes place between segments which agree with respect to some other feature.

3.4.1 The Opacity of Icy Targets

In Khalkha Mongolian (Kaun, 1995; Svantesson et al., 2005), there are distinct processes of tongue root harmony and rounding harmony; it is rounding harmony which is of present interest. In Khalkha, non-high vowels within a word must agree with respect to rounding. Suffixes alternate to take on the rounding specification of the root; unround suffixes appear when no round vowel is present (118ab), and round suffixes appear when a round vowel is present (118cd).

(118) **Khalkha Rounding Harmony**

a. ačaː-gaːr 'burden–INST'

b. deːl-eːr 'coat–INST'

c. nɔxɔi-gɔːr 'dog–INST'

d. doroː-goːr 'stirrup–INST'

High vowels, on the other hand, do not undergo harmony — but the Khalkha inventory, given in (119) below, contains both high and non-high round vowels. In Chapter 2, the non-participation of high vowels is attributed to faithfulness asymmetries — the contrast between high rounded vowels is more salient than the contrast between non-high round vowels, so changes in rounding are more costly (more unfaithful) for high vowels than for their non-high counterparts.

(119) **Khalkha Mongolian Inventory**

i				u
	(ɪ)		ʊ	
e				o
			ɔ	
		a		

High unrounded vowels [i,iː] are transparent; in (120b), for example, the final vowel of the reflexive genitive suffix surfaces as round, despite the unround [iː] intervening between it and the round stem vowel.

(120) **Khalkha Transparency**

a. očidor	'yesterday'	*očider
b. xɔt-iːxɔː	'town–REFL.GEN'	*xɔt-iːxaː
c. nɔir-iːxɔː	'sleep–REFL.GEN'	*nɔir-iːxaː
d. tomr-iːxoː	'iron–REFL.GEN'	*tomr-iːxeː

Despite being round itself, [u,u:] is opaque to harmony. In (121d) and (121h), for example, the perfective suffix surfaces with an unround vowel when the [u:] of the causative intervenes, despite undergoing harmony when there is no intervenor (121bg).

(121) **Khalkha Opacity**

a. ɔr	'enter'	
b. ɔr-ɔːd	'enter–PERF'	
c. ɔr-ʊːl	'enter–CAUS'	
d. ɔr-ʊːl-aːd	'enter–CAUS–PERF'	*ɔr-ʊːl-ɔːd
e. tor	'be born'	
f. tor-oːd	'be born–PERF'	
g. tor-uːl	'be born–CAUS'	
h. tor-uːl-eːd	'be born–CAUS–PERF'	*tor-uːl-oːd

Furthermore, in addition to blocking harmony, high round vowels in initial position fail to initiate harmony on subsequent vowels — the suffixes in the examples in (122) surface as unround, despite being preceded by round vowels, just in case those round vowels are high (cf. 118). The failure of high vowels to trigger rounding harmony can be attributed to the same trigger strength asymmetries that weakened high triggers in Yakut (section 3.3.1 above).

(122) **Khalkha Trigger Failure**

a. uz-leː ‘see–NARR.PAST’ *uz-loː

b. tʊːlai-gaːr ‘hare–INST’ tʊːlɔi-gɔːr

c. guzeː-geːr ‘rumen-INST’ guzoː-goːr

It is the combination of high vowels’ inability to trigger rounding harmony in Khalkha and the fact that high round vowels already bear the spreading feature that leads to their opacity. Because high round vowels are already round, they are able to vacuously link to the spreading feature, satisfying SPREAD(+ROUND) without violating the faithfulness constraint that leads to the absence of harmonic alternations among high vowels. Because high vowels are unable to trigger harmony, however, these segments are unable to subsequently propagate harmony. Though their participation is vacuous, they are **icy targets** (Jurgec, 2011) — segments which undergo harmony but do not trigger it.

In order for preferential triggering (in particular, harmony from preffered triggers *only*) and transparency to coexist in the same language, the weighting condition in (123) must obtain. A reward on SPREAD(F)that is scaled up for preferential triggering and scaled down for non-locality must be sufficient to overcome faithfulness, while an unmodified reward is not. This means that the effect of scaling for trigger strength must be proportionally greater than the effect of scaling for non-locality. If x is 2 and k is 0.5, the result of multiplying by both will be equal to multiplying by neither; however, if x is 2 and k is 0.75, the result of multiplying by both will be greater than multiplying by neither. This indicates that the numerical values of k and x must be able to vary on a language-particular basis; see the analysis of gradient effects in Finnish in Chapter 5 for additional reasons for language-particular differences in the numerical values of the scaling factors.

(123) $w(\text{SPREAD(F)}) * k * x > w(\text{DEP(LINK)}) > w(\text{SPREAD(F)})$

The analysis of transparency, shown in (124), is essentially the same as in Wolof, Hungarian, Yoruba, and others — but this time it’s faithfulness to the rounding contrast in high vowels that prevents them from undergoing harmony, rather than feature co-occurrence. The constant k in the scaling factor for non-locality has been set at 0.75 here to accomodate the weighting conditions in (123). The faithful candidate (124a) receives no violations and no rewards; the candidate with local spreading of [+ROUND] incurs a violation of high-weighted IDENT(ROUND)$_{+\text{HI}}$. The opaque candidate (124c) receives a reward for spreading the non-dominant feature locally. The transparent candidate (124d) receives a reward for spreading the dominant feature — the reward is diminished for non-locality, but increased for spreading from a preferred (non-high) trigger, and the resulting harmony score is still greater than the opaque candidate's. The transparent candidate wins, and the derivation will converge at the second step.

(124) **Step 1**

[+] [−] [−] xɔt-iːxaː	ID(RD)$_{+\text{HI}}$ 5	SPR(+RD) 4	SPR(−RD) 3	DEP(LINK) 3	ID(RD) 2	$\mathcal{H}$
a. [+] [−] [−] xɔt-iːxaː						0.0
b. [+] [−] [−] xɔt-uːxaː	-1	+2 (+1 * 2)		−1	−1	-2.0
c. [+] [−] [−] xɔt-iːxaː			+1	−1		0.0
d. ☞ [+] [−] [−] xɔt-iːxɔː		+1.5 (+1*0.75*2)		−1	−1	1.0

Step 2: *Convergence*

When the medial vowel is itself round, the relationship among the candidates is changed — this is demonstrated in the tableau in (125). At the first step, as above, the faithful candidate (125.1a) earns no rewards and incurs no violations. However, this time spreading the dominant feature locally — from the initial syllable to the medial syllable, in candidate (125.1b) — does not involve a feature change, and is instead an instance of vacuous linking. This means that the high-weighted constraint IDENT(ROUND)$_{+\text{HI}}$, whose violation prevented harmony from applying to the medial vowel in (124) above, is no longer an impediment. Candidate (125.1b) earns a reward on SPREAD(+ROUND) — scaled up because it is induced by a preferred trigger — and its only violation is DEP(LINK).

Candidate (125.1c), which has spread [+ROUND] from the medial high vowel to the final vowel, also earns a reward on SPREAD(+ROUND). However, its reward does not benefit from the scaling factor for preferential triggering, because it is a high vowel and therefore a poor trigger. Candidate (125.1c), which has spread [+ROUND] transparently from the initial vowel to the final vowel, skipping the medial vowel, also earns a reward on SPREAD(+ROUND), and its reward is increased by the scaling factor for preferential triggering, but it is also decreased by the scaling factor for non-locality. Both candidate (125.1c) and candidate (125.1d) also violate IDENT(BACK), because targeting the final vowel involves changing its feature value. Unlike these candidates, candidate (125.1b) benefits from trigger strength and is not impeded by non-locality — and because its vacuous linking doesn't violate the faithfulness constraint that blocked harmony from applying to [i], it is chosen as the winner.

At the second step, the choice is between the faithful candidate (125.2a) and a candidate which has spread [+ROUND] from the medial high vowel to the final vowel (125.2b). Both candidates inherit the reward earned for spreading to the medial vowel, and candidate (125.2b) earns an additional reward for spreading to the final vowel. However, because the trigger is now high [uː], this reward is not increased by the scaling factor for preferential triggering, and is unable to overcome the faithfulness violations required for spreading — reflecting the overall fact that high vowels are not suitable triggers in Khalkha in general. The faithful candidate is chosen as the winner, and the derivation converges immediately.

(125) **Step 1**

[+] [+] [–] tor-uːl-eːd	5 Id(rd)$_{+hi}$	4 Spr(+rd)	3 Spr(–rd)	3 Dep(link)	2 Id(rd)	$\mathcal{H}$
a. [+] [+] [–] tor-uːl-eːd						0.0
b. ☞ [+] [+] [–] tor-uːl-eːd		+2 (+1 * 2)		–1		5.0
c. [+] [+] [–] tor-uːl-oːd		+1		–1	–1	–1.0
d. [+] [+] [–] tor-uːl-oːd		+1.5 (+1*0.75*2)		–1	–1	1.0

Step 2: *Convergence*

[+] [+] [–] tor-uːl-eːd	5 Id(rd)$_{+hi}$	4 Spr(+rd)	3 Spr(–rd)	3 Dep(link)	2 Id(rd)	$\mathcal{H}$
a. ☞ [+] [+] [–] tor-uːl-eːd		+2 (+1 * 2)				8.0
b. [+] [+] [–] tor-uːl-oːd		+3 (+1*2+1)		–1	–1	7.0

In Khalkha, the opacity of high rounded vowels in rounding harmony results from the confluence of two independent properties of the language. The first is that it is greater faithfulness to rounding in high vowels which prevents their participation, rather than feature co-occurrence, as evidenced by the fact that there are both rounded and unrounded high vowels. This means that an already-round high vowel can (and must) be linked to as a dependent, because doing so does not alter its rounding specification, but an unrounded high vowel is unable to do so. Secondly, high vowels are not successful triggers of harmony in general — this was illustrated in the examples in (122), where an initial high rounded vowel failed to spread rounding to a following eligible target. As a result, high rounded vowels vacuously undergo harmony, but fail to propagate it — high unrounded vowels are unable to undergo harmony, and so are skipped.

3.4.2 Parasitic Harmony

Sensitivity to similarity — this time sensitivity between target and trigger with respect to a feature other than the one spreading — can also be seen in parasitic harmony. In these cases, harmony can only occur between segments which already agree along some other feature dimension. In Yawelmani (Archangeli, 1984; Cole and Trigo, 1988; Cole and Kisseberth, 1995b; Kaun, 1995), for example, rounding harmony occurs between vowels of the same height (126) — the suffixes [-hin]∼[-hun] and [-al]∼[-ol] alternate depending on the rounding specification of the root.

(126) **Yawelmani Rounding Harmony**

a. xil-hin 'tangles NON-FUT'
b. dub-hun 'leads by the hand NON-FUT'
c. xat-al 'might eat'
d. bok'-ol 'might find'

In the examples in (126), the root vowel and the suffix vowel are of the same height. If those vowels are of different heights, however, harmony does not occur — in (127b) and (127d), stem and suffix disagree with respect to rounding.

(127) **Yawelmani Parasitic Blocking**

a. xat-hin 'eats NON-FUT'
b. bok'-hin 'finds NON-FUT'
c. xil-al 'might tangle'
d. dub-al 'might lead by the hand'

In Yawelmani, both high vowels (126b) and non-high vowels (126d) are able to trigger and undergo harmony, so the disharmony in (127) cannot be attributed to either restrictions on potential undergoers or restrictions on potential triggers. Harmony in this case is conditioned upon the target and trigger sharing some property, namely height. Other cases of rounding harmony contingent upon height agreement are discussed in Kaun (1995), and examples of ATR harmony contingent upon height agreement include Vata (Kaye, 1982) and Phuthi (Donnelly, 2000).

Parasitic restrictions, like preferential triggering, are perceptually motivated — harmony is more perceptually advantageous when it occurs among otherwise-similar segments (see Chapter 4 for evidence and discussion). The restriction can be accounted for by an additional scaling factor, specific to the relationship between trigger and target, given in (128).

(128) **Scaling factor: parasitic harmony**

For a trigger α, a target β, and a feature F, multiply the reward earned for the dependent segment β by a constant y (such that $y > 1$) if α and β agree along a secondary feature dimension X.

The scaling factor for parasitic harmony exerts its effect in much the same way as the scaling factor for preferential triggering; the weighting conditions are shown in (129). When the weights of Spread(F) and Ident(F) are such that an unmodified reward for harmony is unable to overcome the faithfulness cost, but an increased reward for parasitic harmony is able to do so, the result will be parasitic harmony as in Yawelmani.

(129) $w(\textsc{Spread}(\mathrm{F})) * y > w(\textsc{Ident}(\mathrm{F})) > w(\textsc{Spread}(\mathrm{F}))$

The basic analysis of the Yawelmani data is demonstrated in the tableaux in (130) and (131) below. In (130), the stem and suffix vowels disagree with respect to height. The reward earned for harmony is unscaled, and therefore candidate (130b), which also violates Ident(round), peforms worse than its faithful competitor (130a), which earns no rewards but also incurs no faithfulness violations. The faithful candidate is chosen as the winner, and the derivation converges immediately.

(130) **Step 1**: *Convergence*

[+] [−] bok'-hin	IDENT(RD) 3	SPREAD(+RD) 2	$\mathcal{H}$
a. ☞ [+] [−] bok'-hin			0.0
b. [+] [−] bok'-hun	–1	+1	–1.0

In (131), the stem and suffix vowels agree with respect to height. The faithful candidate (131a) earns no rewards on SPREAD(+ROUND), but also incurs no faithfulness violations. Candidate (131b) incurs a violation of IDENT(ROUND), but also earns a reward on SPREAD(+ROUND). This time, however, because target and trigger agree with respect to height, the reward is multiplied by y (here arbitrarily set at 2). This increased reward is sufficient to overcome the cost of unfaithfulness, and the candidate with harmony is chosen; the derivation will converge at the second step.

(131) **Step 1**

[+] [−] dub-hin	IDENT(RD) 3	SPREAD(+RD) 2	$\mathcal{H}$
a. [+] [−] dub-hin			0.0
b. ☞ [+] [−] dub-hun	–1	+2 (+1 * 2)	1.0

Step 2: *Convergence*

The brief demonstration above abstracts away from the effects of preferential triggering and locality, but these can interact with parasitic restrictions. For example, an alternative analysis of the Yakut example in 3.3.1 above is that it is preference for parasitic harmony that preferential non-high triggers are sufficiently strong to overcome, rather than a bias against the marked vowels that would result from targeting non-high vowels for rounding harmony (Kaun, 1995). Additionally, an alternative analysis of the non-participation (either transparency or opacity) of both high and low vowels in ATR harmony is one in which mid vowels target only other mid vowels precisely because they are both mid (Cole and Trigo, 1988). In both cases, a parasitic harmony analysis is not the only available analysis, but they are both attested examples of the kinds of patterns predicted when parasitic restrictions interact with preferential triggering and locality.

3.5 The Merits of Operational Scaling

While the conceptual framework advocated in this chapter — interpreting the choice betwen transparency and opacity as a competition between local and non-local triggers, subject to influence from other factors determining the strength of each trigger — is essentially the same as that proposed by Hayes and Londe (2006); Hayes et al. (2009), the formal mechanisms by which this conceptual framework is realized differ considerably.

Hayes et al. use a family of Agree-type constraints — Agree(back,local) demands that vowels agree in backness with an immediately preceding back vowel, while Agree(back,non-local) demands that vowels agree in backness with a preceding back vowel anywhere in the word. There are specific versions of each of these constraints which demand agreement with back vowels of height X or lower, and a parallel series demanding that vowels agree in backness with preceding front vowels.

The differences between the present approach and the Hayes et al. approach are too numerous for a direct comparison to be particularly useful — their constraints are negatively-formulated, and situated in a parallel version of OT/HG. Hayes et al. do note the pitfalls of Agree-type constraints, but adopt them for expository simplification to (successfully) illustrate a broader point. Chapter 2 presented arguments in favor of adopting a positive spreading constraint and situating the analysis in a serial framework.

The main point of difference between the approach taken by Hayes et al. (2009) and the present analysis is that they use a family of constraints in stringency relationships, which evaluate agreement between segments as a surface property of representations, while I propose a scalar constraint whose rewards are increased or decreased according to properties of the spreading operation itself. In this section, I will address the specific advantages of scaling rewards in this way — in particular, the ability to capture the generalization that conditions on targets and triggers only ever result in failure of harmony, and do not themselves induce alternations. This will be illustrated with a set of constraints which differs from the present approach only with respect to this relevant difference; in other words, a family of stringent Spread(F) constraints, whose trigger and target conditions are evaluated based on surface representations.

While the harmony constraint itself assigns rewards based on representational properties of candidates — segments which are linked as dependents to some autosegmental feature — the scaling of those rewards is based on the particular instance of spreading which resulted in that segment's status as a dependent (cf. Jurgec 2011, who models local trigger effects via durable representational relationships between triggers and their targets).

The primary benefit of defining triggers and targets operationally is that they are not subsequently altered by changes to the representation. For example, if the trigger/target relationship was a durable property of the representation, inserting material between formerly-adjacent trigger and target would render them non-local — this would wrongly predict languages where epenthesis is blocked just in case it would result in an interuption in the harmony domain. Similarly, removing material between formerly-adjacent trigger and target would increase their locality — predicting a language where vowels shorten `iff` they have been skipped by harmony.

A durable trigger/target relationship would also wrongly predict that, subsequent to spreading, a poor trigger could be compelled to change along another feature dimension to become a better trigger. This kind of behavior is seen in re-pairing harmony *targets*, but their status as preferred or dispreferred targets results not from conditions specific to harmony but from general markedness constraints whose effects are observed elsewhere in the language.

The conditions under which a preference for trigger/target locality could induce shortening of intervening material are illustrated in the tableaux in (132). SPREAD(+ROUND)[d≤V] assigns a reward for a dependent segment if the trigger and target are separated by a distance less than or equal to a short vowel, and SPREAD(+ROUND)[d≤V:] assigns a reward if trigger and target are separated by a distance less than or equal to a long vowel; the constraints are in a stringency relationship. The hypothetical language represented is Khalkha′ — [i] is a non-participant in rounding harmony by virtue of faithfulness asymmetries, and transparency is consistently chosen. Candidates with harmony targeting [i] and candidates with opaque harmony triggered by [i] are assumed to lose for the same reasons they lose in Khalkha, and are not shown. Similarly, the tableaux in (132) abstract away from the scaling effect for perceptually impoverished triggers, since in all candidates the trigger is of the same height.

At the first step of the derivation, the faithful candidate (132.1a) incurs no violations and receives no rewards, for a harmony score of 0. The transparent candidate (132.1b) incurs a violation of IDENT(ROUND), but earns a reward for spreading — because trigger and target are separated by a long vowel, it is SPREAD(+ROUND)[d≤V:] that rewards this candidate. The transparent candidate has a harmony score of 1, and is chosen as the winner.

At the second step of the derivation, the faithful candidate (132.2a) again receives a reward on SPREAD(+ROUND)[d≤V:], and incurs no faithfulness violations, for a harmony score of 3. Its competitor (132.2b) has shortened the long vowel intervening between the trigger and target, violating IDENT(LENGTH). This candidate still earns a reward on SPREAD(+ROUND)[d≤V:] (a short vowel is less than or equal to a long vowel), but now the reduced distance between trigger and target *also* qualifies this candidate for a reward on SPREAD(+ROUND)[d≤V]. The resulting harmony score is 5, and candidate (132b) is chosen as the winner.

(132) *Straw man: distance reduction*

Step 1

[+] [−] [−] pokti:-ra	SPREAD (+RD)[d≤V] 4	SPREAD (+RD)[d≤V:] 3	ID(RD) 2	ID(LENGTH) 2	$\mathcal{H}$
a. [+] [−] [−] pokti:-ra					0.0
b. ☞ [+] [−] [−] pokti:-ro		+1	−1		1.0

Step 2

[+] [−] [−] pokti:-ro	SPREAD (+RD)[d≤V] 4	SPREAD (+RD)[d≤V:] 3	ID(RD) 2	ID(LENGTH) 2	$\mathcal{H}$
a. [+] [−] [−] pokti:-ro		+1			3.0
b. ☞ [+] [−] [−] pokti-ro	+1	+1		−1	5.0

Step 3: *Convergence*

The force compelling vowel shortening in (132) is the pressure to bring the trigger and target closer together — decreasing the distance between trigger and target earns a reward on a higher-weighted version of the harmony constraint, and this outweighs faithfulness to length. Assuming that IDENT(LENGTH) outweighs general markedness constraints against long vowels, Khalkha′ is a language where vowels contrast for length *except when they are transparent*, in which case they and neutralize to short. Vowels like the stem-final [i:] in hypothetical [pokti:] would then alternate for length, depending on whether or not they were followed by a suitable target for harmony.

Similarly, conditions under which a poor trigger could be compelled to change into a better trigger are illustrated in the tableaux in (133). SPREAD(+ROUND)[t=–HI] is a version of SPREAD(+ROUND) which specifically assigns rewards to segments whose (durable) trigger is non-high; SPREAD(+ROUND) applies to all triggers equally (the constraints are in a stringency relationship). This hypothetical language, Khalkha″, is one where high vowels are able to trigger rounding harmony, but non-high triggers are preferred.

At the first step in (133), the faithful candidate (133.1a) incurs no rewards for spreading and no violations of faithfulness; its competitor (133.1b) has spread [+ROUND] from the initial high vowel; it earns a reward on the general SPREAD(+ROUND) constraint, but because spreading is from a high vowel, it does not earn any rewards on the specific version of the constraint rewarding spreading from non-high vowels. Nonetheless, the reward for spreading is still sufficient to overcome the cost of faithfulness, and candidate (133.1b) is the winner.

At the second step, the faithful candidate (133.2a) again earns a reward on the general harmony constraint, but not on the specific version — its trigger is still a high vowel. Its competitor (133.2b) has lowered that initial vowel — because the trigger/target relationship is durable, this affects this candidate's rewards for harmony; now it earns rewards both on the general constraint and on the version specific to non-high triggers. The reward on the higher-weighted constraint that this candidate earns for changing to a preferred trigger easily overcomes the faithfulness cost of changing height, and candidate (133.2b) is the winner.

(133) *Straw man: trigger improvement*

Step 1

[+] [−] puk-ra	SPREAD (+RD)[t=−HI] 4	SPREAD(+RD) 3	ID(RD) 2	ID(HI) 2	$\mathcal{H}$
a. [+] [−] puk-ra					0.0
b. ☞ [+] [−] puk-ro		+1	−1		1.0

Step 2

[+] [−] puk-ro	SPREAD (+RD)[t=−HI] 4	SPREAD(+RD) 3	ID(RD) 2	ID(HI) 2	$\mathcal{H}$
a. [+] [−] puk-ro		+1			3.0
b. ☞ [+] [−] pok-ro	+1	+1		−1	5.0

Step 3: *Convergence*

The language in (133), Khalkha″, is a language where high vowels are successful at triggering rounding harmony, but where a high vowel will lower `iff` it has induced harmony on a following vowel. Scaling rewards based on the harmony operation rather than on durable trigger/target relationships avoids this prediction, and the prediction of Khalkha′, because the reward for harmony is not reassessed unless a new spreading operation is performed.

This is demonstrated in the tableaux in (134). At the first step, the faithful candidate (134.1a) again incurs no rewards and no violations; its competitor (134.1b) earns a reward for spreading [+ROUND], but because the trigger of this operation is a high vowel, the scaling factor for preferential triggering is not applied. Regardless, the reward for spreading is sufficient to overcome the faithfulness cost, and candidate (134.1b) is the winner.

At the second step, both candidates inherit their reward on SPREAD(+ROUND). The faithful candidate (134.2a) incurs no additional violations. Its competitor (134.2b) lowers the initial vowel — but because no spreading operation has taken place, that initial vowel is not a trigger any longer, and so its status as a low vowel confers no benefit. The scaling factors applied to the reward for the dependent segment in (134.2) are based on information about what the trigger was when the linking operation took place — because the trigger was high, the reward remains unscaled, and IDENT(HI) is gratuitously violated. The faithful candidate is chosen as the winner, and the derivation converges.

(134) **Step 1**

[+] [−] puk-ra	SPREAD(+RD) 3	IDENT(RD) 2	IDENT(HI) 2	$\mathcal{H}$
a. [+] [−] puk-ra				0.0
b. ☞ [+] [−] puk-ro	+1	−1		1.0

Step 2: *Convergence*

[+] [−] puk-ro	SPREAD(+RD) 3	IDENT(RD) 2	IDENT(HI) 2	$\mathcal{H}$
a. ☞ [+] [−] puk-ro	+1			3.0
b. [+] [−] pok-ro	+1		−1	1.0

The broader generalization that the operational nature of the scaling factor captures is this: *The only solution to trigger asymmetries is failure of harmony.* Segments don't respond to quality-based triggering asymmetries by changing quality to become better triggers, and segments don't respond to the preference for local triggers by deleting or reducing material to render the trigger and target closer together. The only role that conditions on harmony triggers can play is to prevent harmony from taking place. Because the present theory assesses scaling factors only at the instant of harmony, the influence of influence of those factors is localized derivationally, and their ability to induce other changes is successfully restricted — conditions on harmony are conditions on *operations*, not representations.

3.6 Conclusion

In this chapter, I have argued for a model of transparency and opacity in harmony in which the choice is a choice between competing triggers — a dominant, non-local trigger on the one hand, and a non-dominant local trigger on the other. Factors which influence the strength of each of the competing triggers — in particular, a preference for perceptually impoverished segments as triggers — can influence the choice between transparency and opacity. Good triggers are more likely to be opaque, initiating their own local harmony domain, while poor triggers are more likely to be transparent.

I proposed that this generalization should be encoded as a series of scaling factors on the harmony constraint — rewards are decreased for increasingly non-local triggers, but increased for preferred perceptually impoverished triggers. I demonstrated that the interaction of these scaling factors with each other, with feature dominance, and with faithfulness asymmetries can account for the range of behavior seen in languages in which some non-participating segments are treated as transparent, while others are treated as opaque. Furthermore, I argued that defining the scaling factors over operations rather than representations sufficiently restricts their power, and prevents preferences for locality and for strong triggers from inducing changes other than the prevention of harmony.

Locality and quality-based trigger preferences are not the only factors influencing a segment's ability to act as a trigger for a particular target; I proposed that cases of parasitic harmony, where target and trigger must agree with respect to a feature other than the spreading feature, can also be accounted for with a scaling factor. Finally, the issue of directionality in harmony has been set aside for the present discussion — however, Mullin (2010) provides evidence that directional restrictions behave in many ways as though they are in fact influences on a trigger's strength (for example, overcoming blocking in one direction, but not the other). Directionality, too, could be modeled as an additional scaling factor, permitting a robust account of a broad range of asymmetries seen in harmony processes.

CHAPTER 4

GROUNDING NON-LOCALITY

4.1 Introduction

Central to the discussion of transparency and opacity in harmony is the question of phonetic grounding. One common assumption is that non-local vowel harmony is phonetically unmotivated, and must therefore be derived through a series of independent processes, rather than from a single natural phonological operation. Permitting Gen to create candidates with gapped structures — harmony among non-adjacent segments — and positing a constraint that can exert a preference for such candidates appears to run counter to this assumption.

In this chapter, I address the issue of phonetic grounding and non-local harmony, challenging the assumption that gapped representations are ungrounded. In Section 4.2, I discuss how a purely articulatory theory of the phonetic motivations for harmony depends crucially on strict locality, and cannot account for transparent harmony while maintaining phonetic grounding. I argue that there is also a perceptual motivation for harmony; mulitple manifestations of a phonological feature enhance the perceptual salience of its cues. This perceptual motivation should apply to non-local as well as local harmony.

In Section 4.3 and 4.4, I present experimental evidence in support of the hypothesis that harmony is perceptually motivated. Section 4.3 presents the results of a discrimination study, showing that listeners were faster and more accurate to distinguish nonce words which differed by the features of multiple vowels. Both harmony and dissimilation ensure systems where words will differ by multiple segments; the specific advantage of harmony is at its greatest when segments also agree with respect to some other feature, providing motivation for parasitic restrictions in harmony systems.

Section 4.4 presents the results of a phoneme recall study, showing that listeners were faster and more accurate to identify whether they had heard a target vowel in a preceding nonce word when that word was harmonic for backness/rounding. Crucially, this advantage obtained both in cases of local harmony and in cases of non-local harmony, providing evidence that harmony is perceptually advantageous even when it occurs among non-adjacent segments.

4.2 Grounding harmony

In a functionally grounded theory of phonology, the range of possible phenomena is expected to relate to the set of processes that provide some benefit to either the speaker, the listener, or both. Differing hypotheses about the motivations for harmony have a direct impact on the kinds of harmony processes that we expect to find, and the kinds of harmony processes that a theory of phonology should be expected to generate.

In particular, assuming that harmony is motivated solely on articulatory grounds — as a form of gestural economy — predicts that harmony must necessarily be strictly adjacent, and requires some additional explanation for the existence of transparent harmony. On the other hand, if harmony is also perceptually motivated — as a form of contrast enhancement — limiting the process to strictly adjacent elements is not necessary, and transparent harmony can be included in the set of grounded phonological phenomena.

4.2.1 Articulatory grounding and strict locality

One of the arguments in favor of strict locality centers around the phonetic grounding of harmony. In particular, it has often been assumed that harmony is articulatorily motivated — it serves as a form of effort reduction, since spreading a feature across multiple segments reduces the number of transitions between gestures. This assumption has been implicit in much of the literature on vowel harmony (Archangeli and Pulleyblank, 1994; Bakovic, 2000; Gafos, 1999; Ni Chiosain and Padgett, 2001, and others), but for an explicit analysis based on articulatory ease see e.g. Boersma (1998b); Riggle (1999).

Consider, for example, front/back harmony in Hungarian; a disharmonic word like *[hɔd-nɛk] would require the tongue to transition from back to front in executing its vowels; however, a harmonic word like [hɔd-nɔk] permits a single backing gesture to apply throughout the domain, minimizing gestural transitions. If minimization of gestural transitions is the only grounded motivation for harmony, strict locality is necessary — only harmony among gesturally adjacent segments will provide any benefit. This means that, in the presence of non-participating vowels, opacity should always be preferred over transparency, because it involves fewer transitions. In Hungarian, opaque *[pɔpiːr-nɛk] is predicted, since it involves only a single transition from back to front — transparent [pɔpiːr-nɔk] requires two transitions, from back to front for the transparent vowel and then from front to back again for the suffix vowel. In other words, not only is harmony across a transparent vowel unmotivated, it in fact performs *worse*.

Under this hypothesis about the functional motivations underlying harmony, transparent vowels are only possible if they are not actually transparent. In particular, it has been argued that segments which appear to be skipped by harmony still conform articulatorily to the harmony domain, but this participation is simply not contrastive in the language's inventory (Ni Chiosain and Padgett, 2001; Benus and Gafos, 2007, and others). For example, consonants between rounded vowels in Turkish are pronounced with substantial lip rounding (Boyce, 1990), consistent with the production of a single rounding gesture throughout the relevant domain; because secondary labialization is not contrastive for Turkish consonants, the difference between a plain consonant and a consonant produced with lip rounding does not represent a cross-category alternation.[1] If all examples of apparent transparency could be explained in this manner, it would be possible to maintain strict locality and a purely gesturally grounded account of harmony.

4.2.1.1 Kinande

Kinande (Bantu, Democratic Republic of Congo) is an ATR harmony language with an underlying 7-vowel inventory, and either a 9- or 10-vowel surface inventory (Archangeli and Pulleyblank, 1994; Gick et al., 2006; Mutaka, 1995, and others); the underlying inventory is given in (135).

(135) **Kinande: Underlying Inventory**

i			u
	ɪ	ʊ	
	ɛ	ɔ	
		a	

The competing hypotheses about the surface inventory are given in (136). At issue in Kinande is whether the low vowel [a], which has been described as transparent to ATR harmony, in fact has a [+ATR] counterpart [ʌ]. If [a] does in fact alternate with [ʌ], it is not transparent to harmony, but rather a full participant; harmony, then, would be strictly local.

[1]Boyce argues that this lack of contrast is not by itself sufficient to produce inter-vocalic lip rounding by comparing Turkish to English, where consonantal secondary labiality is similarly non-contrastive, but uCu sequences are coincidental and not the result of harmony. English speakers did not show the same consonant rounding, suggesting that it results from participation in harmony in Turkish (rather than gesture interpolation).

(136) **Kinande: Surface Inventory**

i					u
	ɪ			ʊ	
e					o
	ɛ			ɔ	
		(ʌ)			
				a	

Using both acoustic and articulatory (ultrasound) data, Gick et al. (2006) argue in favor of the latter hypothesis — [a] alternates with [ʌ] in harmonic contexts. The advanced low vowel [ʌ] plays no phonological role in Kinande outside of harmony; alternations between [a] and [ʌ] are non-contrastive. However, Gick et al. argue that the harmonic alternation is in fact a categorical phenomenon, and not simply the result of coarticulation with neighboring vowels. First, they note that the degree of advancement is similar to the degree of advancement shown in other vowels in the inventory which do contrast phonologically for ATR. If the advancement of [a] was simply due to coarticulatory influence from surrounding advanced vowels, the degree of advancement would be expected to be smaller than in the phonologically-driven alternations. However, Gick et al. found that the articulatory and acoustic distance between [a] and [ʌ] is comparable to the distance between [ɪ] and [i].

Furthermore, they demonstrate that the degree of advancement of [ʌ] does not diminish with distance from the triggering vowel. If the advancement of [a] was an effect of coarticulation, this effect would be expected to diminish with distance. For example, in a word like /kágasu̧/ (a proper name), where an advanced high vowel is preceded by a sequence of two low vowels, the initial [a] should be less advanced than the medial [a] if advancement is attributed to phonetic coarticulation with the following [u]. However, Gick et al. found that both low vowels are advanced ([kʌ́gʌsu]), and there was no significant difference between the initial and medial vowels in terms of degree of advancement.

Both the degree of difference between [a] and [ʌ] and the fact that advancement does not diminish with distance serve as evidence that, rather than being a mere coarticulatory effect, the advancement of the low vowel constitutes a phonological alternation. Instead of behaving as transparent to harmony, [a] is in fact a *non-contrastive participant*; it undergoes harmony, despite its alternation lacking contrastive status in the language.

Svantesson et al. (2005) note a similar set of facts in Mongolian — [i] has been described as being transparent to RTR harmony in the language, but does in fact surface as [ɪ] in RTR domains, and behaves like an RTR vowel with respect to other phonological processes. However, in order to maintain articulatorily-motivated strict locality, *all* cases of transparency must be like Kinande and Mongolian. For every language with segments described as being skipped by a harmony process, there should be evidence that those segments are in fact non-contrastive participants.

4.2.1.2 Finnish and Hungarian

In both Finnish and Hungarian, the front unrounded vowels [i] and [e] have been described as transparent to front/back harmony, and are non-contrastive along the front/back dimension (Hayes and Londe, 2006; Kiparsky, 1981; Ringen and Heinämäki, 1999, and many others). The vowel inventory of Finnish is given in (137a), and the vowel inventory of Hungarian is given in (137b).

(137) a. **Finnish: Vowel Inventory**

i(:)	y(:)	u(:)
e(:)	ø(:)	o(:)
æ(:)		a(:)

b. **Hungarian: Vowel Inventory**

i(:)	y(:)	u(:)
e:	ø(:)	o(:)
ɛ		ɔ(:)

If apparent transparency in Finnish and Hungarian is actually non-contrastive participation, as in Kinande, these neutral vowels should be phonetically affected by their harmonic context. Crucially, though, this phonetic effect should share the properties that Gick et al. identify in Kinande to argue in favor of phonological participation rather than low-level coarticulatory influence. The degree to which [i] and [e] are phonetically influenced by back harmony domains should be comparable to the degree of alternation shown by fully participating vowels, and should not diminish with increasing distance from the harmony trigger.

Acoustic studies of [i] and [e] in Finnish have shown that they are affected by their harmony domain — Gordon (1999) and Välimaa-Blum (1999) both found that [i] and [e] have a lower F_2 in back harmonic contexts (when surrounded by or preceded by back vowels) than in front harmonic contexts. For example, Gordon (1999) found that the [i] in words like *tapithan* 'plugs (emphatic)' had a lower F_2 than the [i] in words like *tätihän* 'aunt (emphatic)'.

Neither study provides control conditions with fully participating vowels for a conclusive comparison, however it is unlikely that the degree to which [i] and [e] are phonetically influenced by harmonic contexts is comparable to that of phonologically participating vowels. Gordon (1999) describes the difference between neutral vowels in back harmonic contexts and in front harmonic contexts as "much smaller" than the equivalent alternations in non-neutral vowels. Furthermore, Välimaa-Blum (1999) provides data from two different speech rates; while the difference between neutral vowels in back contexts and neutral vowels in front contexts was statistically significant in normal speech, it did not reach significance in careful speech.

While further studies are needed to determine how the size of the F_2 difference in neutral vowels in front and back harmonic contexts relates to the size of the F_2 difference in fully participating vowels, the evidence available suggests that Finnish differs from Kinande in this respect. The neutral vowels in Finnish, while affected by the quality of surrounding vowels, do not appear to be affected to a sufficient degree to warrant status as non-contrastive participants.

Benus and Gafos (2007) present the results of an articulatory study of transparent vowels in Hungarian. Using both ultrasound and electromagnetometry, they show that the articulations of [i] and [e] are affected by the quality of surrounding vowels. For example, the [i] in words like *zafír-ban* 'saffire' was articulated with a tongue position that was less front than the [i] in words like *zefír-ben* 'zephyr'. While Benus and Gafos do not provide an explicit comparison between the degree of difference shown by [i] and [e] and the degree of difference shown in fully participating vowels, their data provide indications that the neutral vowels in Hungarian differ from the non-contrastive participants in Kinande.

In Kinande, the similarity in degree of advancement between non-contrastive participants and contrastive participants is consistent with a constant ATR gesture throughout the harmonic domain. The electromagnetic midsagittal articulometry (EMMA) data provided by Benus and Gafos (reproduced here in Figure 4.1), on the other hand, show a clear interruption of the back gesture. It is clear from following the horizontal position of the tongue that both the tongue dorsum and the tongue blade are substantially retracted during the initial [ɔ] in [zɔfirbɔn], and again during the final [ɔ]. However, during the medial [i], the tongue dorsum and tongue blade show substantial fronting, consistent with a categorically front vowel.

If [i] was a non-contrastive participant in Hungarian, as in Kinande, the tongue dorsum and tongue blade would be expected to remain retracted throughout the entire domain. The fact that there is substantial advancement during the medial

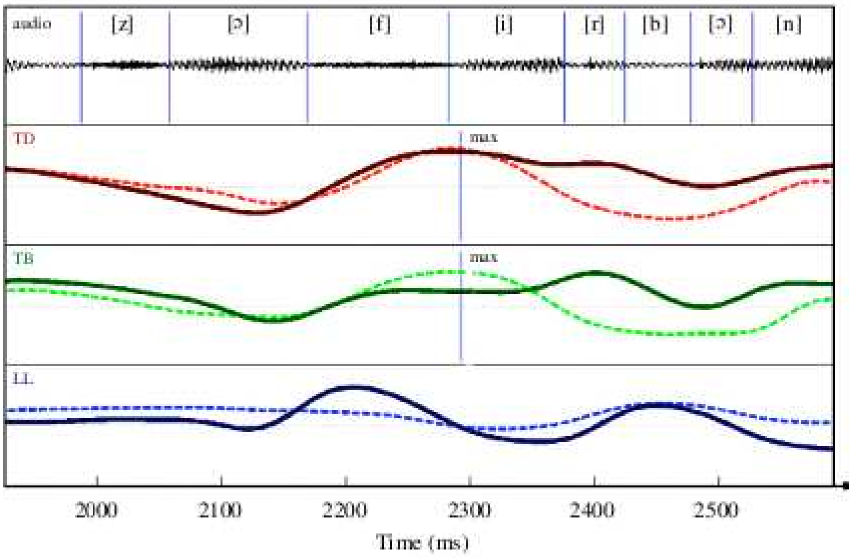

Figure 4.1. From Benus and Gafos (2007, p. 278), EMMA data for the articulation of *zafirban*. Vertical (solid line) and horizontal (dashed line) location of the tongue blade (TB), tongue dorsum (TD), and lower lip (LL).

neutral vowel suggests that [i] is not in fact a participant in the harmony domain. Furthermore, even if t could be considered to be a participant at the phonological level, the fact remains that the implementation of a harmony domain with a medial [i] involves an interruption of the back gesture.

If the purpose of harmony is to minimize gestural transitions, transparent harmony in Hungarian has failed to accomplish that objective. A word like *zafirban* involves a transition from back to front into the medial [i], and then a second transition from front to back for the final [ɔ]. The phonologial process at work in Hungarian, then, results in a form which is, in its articulatory execution, incompatible with a gestural motivation for harmony.

4.2.1.3 Phonetic Grounding and Preserving Strict Locality

Transparent harmony in Hungarian does in fact involve multiple gestural transitions. While there have been proposals aiming to account for transparency while maintaining strict locality in the phonology (see for example Benus 2005; Ni Chiosain and Padgett 2001; Walker 1998; see Chapter 6 for further discussion of proposals of this type), the existence of gestural discontinuity in the surface phonetics seriously undermines a central argument for the desirability of such an account. If adherence to strict locality at the phonological level is motivated by the necessity of gestural locality in the surface phonetics, the existence of harmonic alternations involving discontinuous gestures removes the motivation for phonological strict locality.

One premise behind these proposals is that even patterns that appear to lack phonetic grounding on the surface should be able to be decomposed into separate phonetically motivated processes. However, while it is possible to derive non-local harmony with reference to only adjacent elements, it is not true that each of the necessary mechanisms involved maintains phonetic grounding.

To produce the pattern attested in Hungarian, for example, Benus (2005) proposes two separate mechanisms for driving harmony. The first mechanism is Agree(CL), a constraint demanding minimal articulatory distance (along the front/back dimension) between adjacent segments. This constraint is gradiently evaluated, and interaction with faithfulness constraints can produce either categorical harmonic alternations or gradient coarticulation.

The second mechanism involves assessing the degree of retraction of a given segment — the distance between its actual articulatory target and the underlying or canonical articulatory target). A constraint Agree(R) requires that a segment with a sufficient degree of retraction is followed by a segment that is back. Because neutral [i] and [e] undergo coarticulation in harmonic environments, Agree(R) ensures that a coarticulated [i] or [e] will trigger the following segment to surface as back, mimicking the effects of non-local harmony.

While this account preserves strict locality, and is phonetically grounded insofar as the mechanisms used refer directly to properties of articulation, it is not *functionally* grounded; AGREE(R) in fact prefers candidates with greater articulatory distance between segments. The articulatory distance between slightly retracted [i̠] and a front vowel is less than the distance between [i̠] and a back vowel — if minimizing articulatory distance is the motivation for harmony, [i̠] should be followed by front vowels, not back vowels.

Given the existence of harmony patterns with surface non-locality, and given that analyses preserving strict locality at the phonological level still involve departures from its articulatory motivation, there is reason to doubt the widespread assumption that maintaining strict locality provides a more phonetically grounded account of harmony.

4.2.2 Perceptual grounding and non-locality

The economy of gestural transitions is not the only possible phonetic grounding for harmony. A variety of perceptual motivations have been proposed, and these motivations are not restricted to strictly local harmony.

Suomi (1983) proposes that front/back harmony is a strategy for facilitating the perception of F2 contrasts. He notes that languages with front/back harmony are typically languages in which front vowels also contrast for rounding — for example, in Turkish, Finnish, and Hungarian, front vowels may be either rounded or unrounded. Because the back/front and the rounded/unrounded contrasts are both manifested as changes in F2, having vowels that conrast for both backness and rounding requires the listener to perceive F2 fairly precisely in order to correctly identify the quality of the vowel in question.

Suomi proposes that harmony reduces the perceptual burden on the listener by rendering the front/back contrast predictable. Harmony in the languages he discusses is left-to-right, and the front/back value of non-initial vowels can be predicted on the basis of that of the initial vowel. Because initial syllables are psycholinguistically prominent, the likelihood of accurate F2 perception there is higher than in non-initial syllables. Because non-initial syllables are predictable, the listener does not need to carefully attend to subsequent cues for backness. Furthermore, because subsequent vowels agree in backness with the initial vowel, the cues in those vowels can serve to reinforce the initial hypothesis that a listener has made about the feature value of the initial vowel, further increasing the likelihood of an accurate identification.

Kaun (1995), building on Suomi's proposal, focuses on the hypothesis that extending the duration of realization of a particular feature enhances its perceptibility. She notes, following Steriade (1995), that perceptually difficult contrasts that are subject to positional neutralization tend to be licensed in positions of greater duration (e.g. word edges and prosodically strong positions). This permits not only a more complete articulatory realization, but also a more robust opportunity for the listener to correctly identify the cue.

In vowel harmony, rather than simply limiting a difficult contrast to a position of greater duration, the cues are realized across multiple segments — extending the duration of a feature's realization and thus increasing the likelihood of a correct identification. Like Suomi, Kaun proposes that in harmony domains, cues to the quality of later vowels serve to reinforce a listener's initial hypothesis about the feature value in question.

Gallagher (2010) discusses the perceptual advantages of harmony not in terms of feature identification but in terms of maximizing the distinctness of lexical items. She argues that assimmilatory co-occurrence restrictions on laryngeal features are driven by the pressure for roots to be as perceptually discriminable as possible. In a language with laryngeal assimilation, words with different specificaitons for laryngeal features consistently differ with respect to those features on *all* relevant consonants. Words with multiple differences are more perceptually distinct from one another than words with only a single difference, so laryngeal assimilation aids in perception.

Gallagher supports this claim experimentally, presenting results of several discrimination studies in which subjects were presented with pairs of disyllabic (CVCV) nonce words and asked to make a same/different judgment. In the first study, each nonce word contained either zero, one, or two ejectives; in the second study, each nonce word contained either zero, one, or two aspirates. Results from both studies patterned similarly — subjects were significantly more accurate at discriminating words that differed by laryngeal features in both consonants than words that differed by laryngeal features in only one. On other words, a contrast between zero and two ejectives/aspirates was more readily perceptible than either a contrast between zero and one or a contrast between one and two ejectives/aspirates.

Gallagher's results are based on the perception of consonantal rather than vocalic features, but the broader effect should be general — in a language with harmony, words are consistently more different from each other than in a language without harmony, and are therefore more easily distinguished.

Suomi (1983), Kaun (1995), and Gallagher (2010) each take a slightly different perspective on the perceptual advantages of harmony, but the central claims are similar: harmony is advantageous because it enhaces the perceptual salience of difficult contrasts. Unlike an articulation-based explantion, the perceptual advantages of harmony do not require strict locality. Even when harmony is non-contiguous, it can serve to render contrasts predictable, provide additional reinforcement for the identification of a feature value, and ensure that words are more different from each other than they would otherwise be. Under this approach, non-local harmony is still phonetically grounded.

This is not to say, of course, that local and non-local harmony are *equally effective* at achieving the motivating perceptual advantages. Many of these benefits rely on the listener interpreting a sequence of vowels bearing the same feature as representing a single contrast — for example, using harmony targets to bolster a hypothesis about the quality of an initial trigger is only possible if those targets are interpreted as additional manifestations of the same feature present on the trigger.

Anything that reduces the likelihood of a listener interpreting multiple cues as representing a single contrast, then, should diminish the effectiveness of harmony as a perceptual strategy. It's reasonable to hypothesize that listeners will be less likely to interpret two segments bearing the same feature as representing a single contrast if those segments are not adjacent. Non-local harmony, therefore, is expected to be less effective than local harmony as a means of increasing the perceptual salience of a contrast.

This source of explanation maintains phonetic grounding in non-local harmony, and reflects the implicational relationship between local and non-local harmony. The difference between the motivation for local harmony and the motivation for non-local harmony is *quantitative* rather than *qualitative* — they are both beneficial for the same reason, but to differing degrees. As a consequence, (less effective) non-local harmony should imply the presence of (more effective) local harmony. This is in fact true of harmony systems generally; any language with non-local (transparent) vowel harmony also has harmony among adjacent vowels, but not vice versa.

Kaun notes that perceptual grounding of harmony also provides an explanation for certain segments' tendency to preferentially trigger harmony. Since harmony serves to enhance perceptually weak cues, the segments that are in most need of spreading their features are those whose cues are particularly impoverished. See Chapter 5 for further discussion of the relationship between perceptual grounding and preferential triggers.

The perceptual account discussed here should not be taken to be the only motivating factor driving harmony. The articulatory pressure to minimize gestural transitions is likely also a factor — local harmony can be seen as satisfying both perceptual and articulatory pressures. The purpose of this chapter is not to argue that the articulatory advantage does not obtain, but simply that it is insufficient to explain the full range of harmony patterns attested. Perceptual grounding, on the other hand, goes some considerable distance towards providing such an explanation.

Boosts in perceptual salience are not the only advantages of harmony which do not necessarily require strict locality. For example, vowel harmony has been shown to aid in speech segmentation. Suomi et al. (1997) and Vroomen et al. (1998) presented Finnish-speaking subjects with sequences of three CV syllables, the latter two of which constituted an actual Finnish word; the task was to press a button when they heard a word, and to say what that word aloud. Subjects were significantly faster to recognize the target word when the stray syllable disagreed with the target in backness (e.g. puhymy) than when it agreed with the target (pyhymy).

Because the domain of application for harmony processes is usually the word, Suomi; Suomi et al. suggests that harmony serves a delimatative function, providing a source of evidence for locating word boundaries in connected speech.

Like the perceptual advantages discussed above, the benefits of harmony as a cue to word boundaries does not depend crucially on strict locality. Suomi et al. and Vroomen et al. do not have any neutral vowels in their materials, but even non-local harmony serves to delimit the boundaries of the word which is its domain. In a transparent sequence like F [B N B] F, the F B and B F strings still serve to indicate boundaries; that second boundary cue is missing, however in an opaque sequence like F [B N F] F, where harmony has not applied non-locally.

Further research is needed to investigate the effects that neutral vowels have on the role of harmony in word segmentation, and that line of inquiry will be set aside for now. The remainder of this chapter provides experimental evidence supporting the proposal that harmony is perceptually grounded, motivated because it enhances the salience of weakly cued contrasts.

4.3 Experiment 1: Discrimination

The purpose of this experiment is to test Gallagher (2010)'s claim that harmony aids perception by making words more different from each other (and hence more easily discriminable), attempting to replicate with vowel contrasts the results she obtained with contrasts for laryngeal features. Furthermore, I test the hypothesis that the plausibility of a single-contrast interpretation affects the degree to which harmony is advantageous — I compare harmony between vowels that agree with respect to a secondary feature with harmony between vowels that do not.

In languages with parasitic harmony, agreement with respect to some (primary) feature is dependent on agreement with respect to another (secondary) feature. For example, in Yawelmani, rounding harmony is parasitic on height (Cole and Kisseberth, 1995b, and others); harmony occurs between stem and suffix vowels that agree in height, but not between vowels of differing height (138). Both high vowels (138b) and non-high vowels (138d) are able to trigger rounding harmony, and both are able to undergo harmony, but a non-high vowel cannot trigger harmony on a high vowel (138f), and a high vowel cannot trigger harmony on a non-high vowel (138h).

(138)

Height Agreement: Rounding Harmony		Height Disagreement: No Harmony	
a. xil-hin	'tangles NON-FUT'	e. xat-hin	'eats NON-FUT'
b. dub-hun	'leads NON-FUT'	f. bok'-hin	'finds NON-FUT'
c. xat-al	'might eat'	g. xil-al	'might tangle'
d. bok'-ol	'might find'	h. dub-al	'might lead'

Rounding harmony is also parasitic on height in a variety of Turkic languages; see Kaun (1995) for a survey. Additionally, ATR harmony may be parasitic on height, as it is in Vata (Kaye, 1982) and Phuthi (Donnelly, 2000).

Kaun explains the motivation for parasitic restrictions in articulatory terms. Citing Linker (1982), she notes that the rounding gesture in high vowels is both quantitatively and qualitatively different than the rounding gesture in non-high vowels. She therefore proposes a constraint, UNIFORMITY, requiring a feature to be implemented with a uniform gesture.

However, an alternative explanation is the hypothesis that the advantage of realizing a harmonizing feature with a consistent, uniform gesture lies in its perceptual consequences. If harmony is perceptually advantageous because realizing a feature across multiple segments aids in identifying that feature's value, then listeners must be interpreting those segments as reflexes of a single feature. This interpretation is more likely if the realization is consistent across segments, as it is when rounding is spreading among vowels that already agree in height. When rounding is spreading among vowels which disagree in height, however, the acoustic effects of that rounding will differ throughout the course of the feature's realization. This should interfere with a listener's interpretation of the harmonized segments as manifesting a single contrast, therefore diminishing the perceptual advantage of harmony.

Under this account, parasitic restrictions represent a strategy for deploying harmony only where it will be the most beneficial — among segments that agree with respect to some secondary feature. While Kaun (1995)'s discussion is limited to rounding, the same reasoning applies to cases where ATR is parasitic on height. The boost in salience achieved by spreading ATR is greatest among vowels which also agree with respect to height.

In this experiment, subjects are presented with pairs of disyllabic nonce words and asked to make a same/different judgment. There are two distinct yet closely related hypotheses under consideration. The first hypothesis is that harmony is perceptually advantageous because it maximizes the discriminability of words — subjects should be faster and more accurate at distinguishing pairs of words where *both* vowels are different with respect to the harmonic feature (as is the case in languages with harmony) than pairs of words where only one vowel differs (as is the case in languages without harmony). The second hypothesis is that this advantage is diminished when harmony occurs among vowels which disagree with respect to a secondary feature — subjects should be faster and more accurate at distinguishing pairs of words where each word is internally harmonic for both the primary and secondary feature than pairs of words where each word is internally harmonic for one feature but not the other.

4.3.1 Methods

4.3.1.1 Stimuli

The primary harmonzing feature used in this experiment is ATR, and the secondary feature is height. The stimuli were disyllabic nonce words, of the shape CVCV; the consonants [h] and [g] were chosen because of their low coarticulatory effect on surrounding vowels, and vowels were manipulated for height (high or mid) and ATR. A list of the resulting nonce words can be found in the left-hand column in Table 4.1.

Each CV syllable was recorded separately, in a neutral frame sentence, read by a phonetically trained native speaker of North American English in a sound-attenuated booth. For [+ATR] vowels, the diphthong portion was removed. Using Praat (Boersma and Weenink, 2008), syllables were equalized for F_0 (220hz), consonant duration (50ms), and vowel duration (150ms), then spliced together to form the nonce words. Based on information from a pilot study, a small amount of pink noise was added.

While there are some differences between the tense/lax contrast used in languages like English and the ATR contrast in languages with tongue root harmony (see e.g. Lindau-Webb 1987), the manipulations above resulted in vowel stimuli that closely resembled the acoustic properties typical of ATR languages (see Starwalt 2008 for further discussion of those acoustic characteristics).

	Stimulus Pairs			
Items	*Same*	*Different (1)*		*Different (2)*
hege	hege--hege	hege--hegɛ	hige--higɛ	hege--hɛgɛ
hegɛ	hegɛ--hegɛ	hege--hɛge	hige--hɪge	hegɛ--hɛge
hɛge	hɛge--hɛge	hegɛ--hege	higɛ--hige	hɛge--hegɛ
hɛgɛ	hɛgɛ--hɛgɛ	hegɛ--hɛgɛ	higɛ--hɪgɛ	hɛgɛ--hege
hegi	hegi--hegi	hɛge--hege	hɪge--hige	hegi--hɛgɪ
hegɪ	hegɪ--hegɪ	hɛge--hɛgɛ	hɪge--hɪgɛ	hegɪ--hɛgi
hɛgi	hɛgi--hɛgi	hɛgɛ--hegɛ	hɪgɛ--higɛ	hɛgi--hegɪ
hɛgɪ	hɛgɪ--hɛgɪ	hɛgɛ--hɛge	hɪgɛ--hɪge	hɛgɪ--hegi
hige	hige--hige	hegi--hegɪ	higi--higɪ	hige--hɪgɛ
higɛ	higɛ--higɛ	hegi--hɛgi	higi--hɪgi	higɛ--hɪge
hɪge	hɪge--hɪge	hegɪ--hegi	higɪ--higi	hɪge--higɛ
hɪgɛ	hɪgɛ--hɪgɛ	hegɪ--hɛgɪ	higɪ--hɪgɪ	hɪgɛ--hige
higi	higi--higi	hɛgi--hegi	hɪgi--higi	higi--hɪgɪ
higɪ	higɪ--higɪ	hɛgi--hɛgɪ	hɪgi--hɪgɪ	higɪ--hɪgi
hɪgi	hɪgi--hɪgi	hɛgɪ--hegɪ	hɪgɪ--higɪ	hɪgi--higɪ
hɪgɪ	hɪgɪ--hɪgɪ	hɛgɪ--hɛgi	hɪgɪ--hɪgi	hɪgɪ--higi

Table 4.1. Pairs of nonce-word stimuli.

4.3.1.2 Subjects and Task

The experiment used an AX discrimination task. Subjects were 36 native speakers of North American English, students in introductory linguistics courses recruited by offering extra credit points. They were presented with the nonce word stimuli in pairs, with an ISI of 500ms between words, and asked to make a same/different judgment. In the pairs, given in Table 4.1, the two nonce words differed with respect to the ATR value of one of the vowels, both of the vowels, or neither of the vowels (the "same" condition).

Stimuli were presented in blocks, with a 96 trials per block. There were 6 blocks, and each subject saw each "same" pair a total of 18 times, and each "different" pair a total of 6 times, resulting in an equal number of trials for which "same" and "different" were the correct responses. Responses were cut off at 1500ms, and no subject failed to respond on more than 10% of the trails.

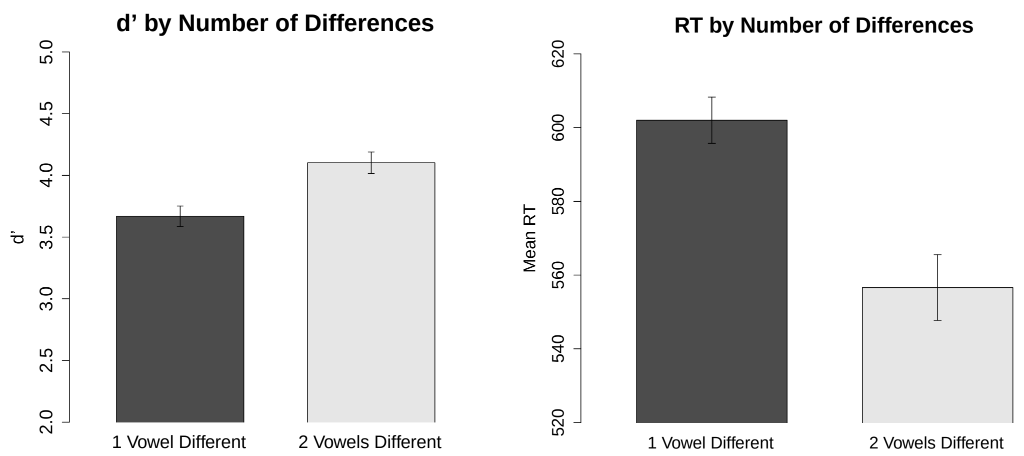

Figure 4.2. Accuracy and response time as a function of the number of differing vowels between members of stimulus pairs, averaged across all subjects and all items. Error bars represent 95% confidence intervals.

4.3.2 Results

4.3.2.1 One Difference vs. Two

As seen in Figure 4.2, subjects were significantly faster and more accurate to distinguish pairs of stimuli where the members of the pair differed by the ATR values of both vowels (e.g. hege ∼ hɛgɛ) than when the members of the pair differed by the ATR value of only one member of the pair (e.g. hege ∼ hegɛ).

For response times, only the "different" pairs were analyzed. A Linear Mixed Effects Model[2] with log-transformed response time as the dependent variable and random effects for subject and item found the difference in response time to be significant ($t = -4.83$, $p < 0.0001$). Accuracy was measured using an independent-response d' measure. A Linear Mixed Effects Model with d' as the dependent variable and ran-

[2]Mixed models were done using the `lme4` package (Bates and Maechler, 2010) in R (R Development Core Team, 2009), and p-values were obtained via Markov Chain Monte Carlo sampling using the `languageR` package (Baayen, 2010). The `lme4` package does not provide an accurate way of estimating degrees of freedom, so those are omitted.

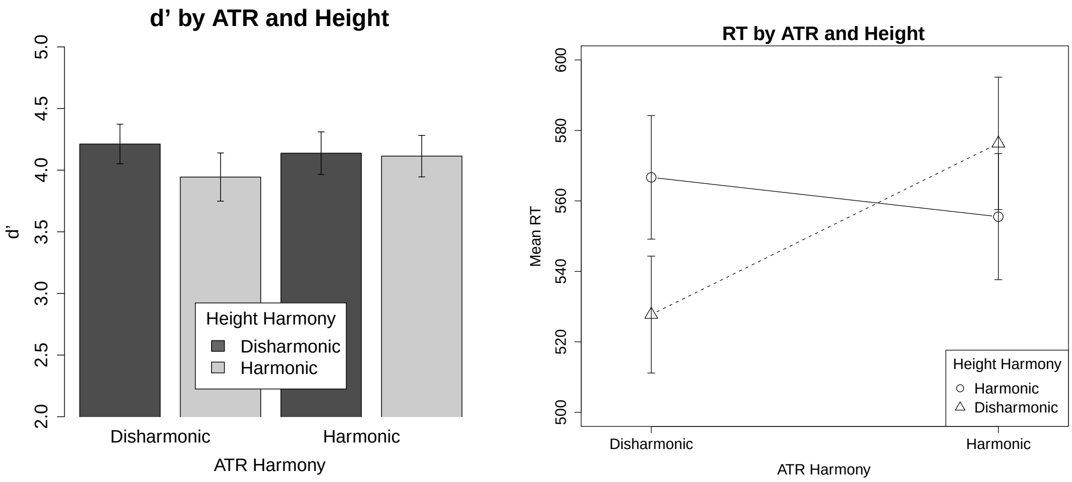

Figure 4.3. Accuracy and response time as a function of height harmony and ATR harmony, averaged across all subjects and all items, for pairs where both vowels differed. Error bars represent 95% confidence intervals.

dom effects for subject and item found the difference in accuracy to be significant (t = 4.46, $p < 0.001$).

4.3.2.2 Height and ATR Interaction

Among pairs where the members of the pair differed by the ATR value of both vowels, there was a significant interaction between ATR harmony and height harmony. For pairs where each nonce word was internally harmonic for height, subjects were slightly (but not significantly) faster to respond to pairs where each nonce word was also internally harmonic for ATR (e.g. hege ∼ hɛgɛ) than pairs where each nonce word was internally disharmonic for ATR (e.g. hegɛ ∼ hɛge). However, for pairs where each nonce word was internally disharmonic for height, subjects were significantly faster to respond to pairs where each nonce word was also internally *disharmonic* for ATR (e.g. hegɪ ∼ hɛgi) than pairs where each nonce word was internally harmonic for ATR (e.g. hegi ∼ hɛgɪ).

Accuracy was high across the board, and a Linear Mixed Effects model with d′ as the dependent variable and random effects for both subject and item found only a significant difference between of ATR harmony within pairs of height-disharmonic words ($t = -2.21$, $p < 0.05$)[3], but in no other conditions. Because accuracy was so high across all conditions, this lack of difference is likely attributable to a ceiling affect.

A Linear Mixed Effects Model with log-transformed response time as the dependent variable and random effects for both subject and item found that, within pairs of height-disharmonic words, there was a significant effect of ATR harmony ($t = 2.61$, $p < 0.01$).[4] Within pairs of height-harmonic words, the effect of ATR harmony failed to reach significance ($t = 0.9$, $p > 0.05$).[5] There was a significant interaction between height harmony and ATR harmony ($t = -2.48$, $p < 0.05$).

In light of this interaction, a comparison was made between height-disharmonic pairs whose members differed by the ATR value of one vowel (e.g. hegi ∼ hɛgi) and height-disharmonic pairs whose members differed by the ATR value of both vowels and were harmonic for ATR (e.g. hegi ∼ hɛgɪ). Figure 4.4 shows that subjects were faster and more accurate in the ATR harmonic pairs differing by two vowels than the pairs differing by only one vowel.

A Linear Mixed Effects Model with log-transformed response time as the dependent variable and random effects for both subject and item found this difference in response time to be significant ($t = -3.95$, $p < 0.001$). A Linear Mixed Effects Model with correct responses as the dependent variable and random effects for both subject and item also found the difference in accuracy to be significant ($t = 3.42$, $p < 0.001$).

[3]Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

[4]Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

[5]Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

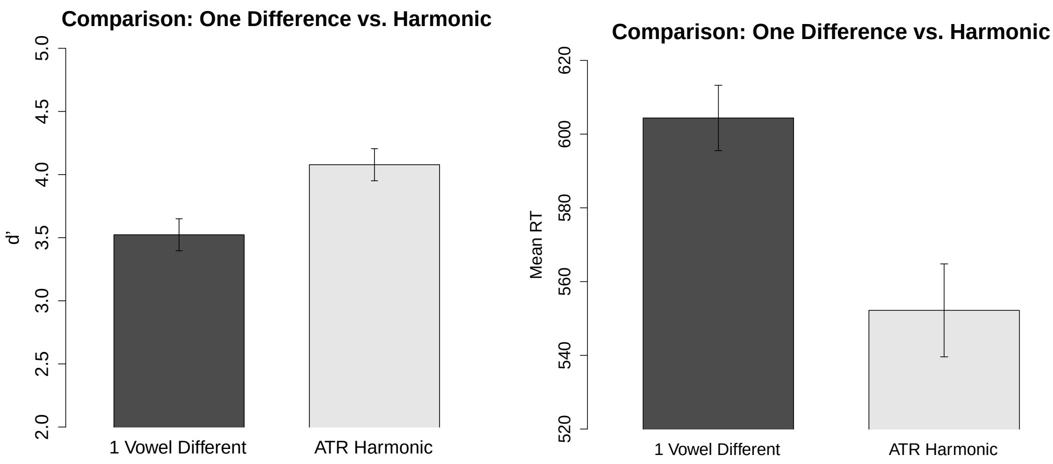

Figure 4.4. Accuracy and response time for pairs of height-disharmonic words, averaged across all subjects and all items, comparing pairs whose members differed by one vowel to pairs whose members differed by both vowels and were ATR harmonic. Error bars represent 95% confidence intervals.

4.3.3 Discussion

The results of this experiment support the claim that harmony is perceptually advantageous because it maximizes the distinctness of words, and that this advantage is diminished in contexts where listeners are less likely to interpret multiple segments bearing the same feature as representing a single contrast.

The central result (shown in Figure 4.2) that words differing by multiple segments are distinguished faster and more accurately replicates one of the results of Gallagher (2010)'s study of laryngeal contrasts, and should hardly come as a surprise. Words differing by multiple segments are more recognizable as different because they are, in fact, more different — indeed, it would be surprising if this *wasn't* the case.

However, the fact that this result is unsurprising should not be mistaken for triviality. Gallagher shows how the pressure to maximize the distinctness of words can be used to explain both assimilatory and dissimilatory co-occurrence restrictions — either type of restriction ensures an overall system where distinct words consistently differ by the feature values of multiple segments.

Of present interest is the specific advantage of *harmony* as a strategy for maximizing perceptual distinctness. The results in Figure 4.3 can be seen as addressing the tradeoff between assimilatory and dissimilatory restrictions. Pairs that differ by two vowels and are ATR harmonic represent forms that obey an assimilatory co-occurrence restriction, and pairs that differ by two vowels and are ATR disharmonic represent forms that obey a dissimilatory co-occurrence restriction.

In Figure 4.3 we see that, when the vowels in a word agree with respect to a secondary feature (height), there is no particular advantage of either assimilation or dissimilation — both are more or less equally effective at easing the perceptual burden to the listener, because both ensure that the words in the pair differ by both vowels. However, when the vowels of a word *disagree* with respect to a secondary feature, harmony is at a clear disadvantage.

Parasitic harmony, then, is a way of deploying harmony only when it will be most perceptually beneficial — when both vowels agree with respect to some secondary feature. Crucially, non-parasitic harmony is still advantageous — in Figure 4.4, we see that even when words disagree with respect to the secondary feature, an assimilatory co-occurrence restriction that ensures multiple differences between words would limit the set of possible comparisons to those which are faster and more accurate (the ATR harmonic pairs) than the comparisons in a language with no co-occurrence restriction (the pairs differing by only one vowel).

The pattern of perceptual advantage seen in this experiment reflects the attested typology of harmony systems — harmony is beneficial even when it is not parasitic, but it is even more beneficial when there is secondary-feature agreement. There are languages which use harmony across both secondary-agreement contexts and secondary-disagreement contexts, and languages which make use of harmony only in secondary-agreement contexts, but crucially there are no languages where harmony is *anti*-parasitic — no languages where vowels assimilate only if they disagree with respect to some other feature.

This does raise a question, though, about the advantages of dissimilatory co-occurrence restrictions for vowels. While vowel harmony is robustly attested, vowel dissimilation is strikingly rare — it occurs in Woleian (Sohn, 1975) and other Oceanic languages (see Lynch 2003 for further discussion). Dissimilation is, however, robustly attested for consonantal features (Alderete, 1997; Gallagher, 2010; Odden, 1994; Yip, 1988, and many others). One explanation for this difference concerns articulation — while we saw in Section 4.2.1 that the gestural motivations of harmony can't account for the full set of facts, the articulatory advantages still obtain. Harmony is advantageous both perceptually and articulatorily (at least when it is local); dissimilation is only advantageous perceptually, and is articulatorily marked.

The prevalence of consonant dissimilation may be due to a difference between consonant articulation and vowel articulation — vowel articulations overlap considerably, while consonant articulations (particularly those for e.g. laryngeal features) overlap substantially less (or not at all). An additional prediction of the results of this experiment is that dissimilation should be anti-parasitic — because dissimilation is of the most benefit in contexts of secondary-feature disagreement, we expect to find languages which dissimilate only in those contexts. While there are too few languages with vowel dissimilation to know if this prediction is borne out, it is attested in consonant co-occurrence restrictions — see Gallagher (2010) for a discussion of languages with "identity effects", where consonants dissimilate for laryngeal features only if they are non-homorganic.

4.3.4 Summary

This section presented the results of a discrimination study designed to test the hypothesis that harmony is perceptually advantageous because it maximizes the perceptual distinctness of words. Furthermore, the experiment tested the hypothesis that this advantage is diminished in cases of secondary-feature disagreement, where listeners would be less likely to interpret multiple segments bearing the same feature as representing a single contrast.

Subjects were indeed faster and more accurate to discriminate words which differed by feature values of multiple segments than words which differed by only a single segment. Furthermore, this advantage was diminished in words with secondary-feature disagreement, supporting the claim that parasitic harmony patterns are perceptually motivated.

Discrimination is not an ideal test for the specific advantages of harmony, because dissimilatory co-occurrence restrictions also ensure a system where words consistently differ by multiple segments. The following section presents the results of an experiment testing the hypothesis that harmony is perceptually grounded using a slightly different methodology.

4.4 Experiment 2: Phoneme recall & non-locality

The purpose of this experiment is to provide support for the claim that harmony is perceptually advantageous because it aids in feature recognition, and to test the hypothesis that this advantage obtains even when the segments bearing the feature in question are non-adjacent.

While the previous experiment tested the perceptual grounding hypothesis using a discrimination task, this experiment uses a phoneme recall task — subjects hear a nonce word, followed by an isolated vowel, and are asked to indicate whether they heard that vowel in that nonce word.

If Kaun (1995)'s hypothesis is correct, it should be easier for subjects to recognize that they heard a vowel with a particular feature value if there was another vowel with that feature value in the word — multiple cues to a feature should aid in its

	Disharmonic		Harmonic	
Local	u e a	e u a	i e a	e i a
	i o a	o i a	u o a	o u a
	i u a	u i a	i i a	u u a
Non-Local	u a e	e a u	i a e	e a i
	i a o	o a i	u a o	o a u
	i a u	u a i	i a i	u a u

Table 4.2. Vowel combinations used in nonce-word stimuli.

recognition. In principle, this advantage should also obtain even when those cues are non-adjacent.

Unlike the discrimination task used in the previous experiment, this task is testing something more specific about harmony in particular. The advantage of harmony in discrimination is also an advantage of dissimilation — in this study, any benefit that harmony provides in feature recognition will be more specifically attributable to harmony.

4.4.1 Methods

4.4.1.1 Stimuli

The features under consideration in this study were backness and rounding (which always covaried, as this experiment was performed on English-speaking listeners).[6] Stimuli were trisyllabic (CVCVCV) nonce words, and isolated vowels. For the nonce words, the consonants [h], [k], and [g] were chosen because of their low coarticulatory effect on surrounding vowels. For the isolated vowels, a [ʔ] onset was used.

Each CV syllable was recorded separately, in a neutral frame sentence, read by a phonetically trained native speaker of North American English in a sound-attenuated booth. Using Praat, syllables were equalized for F_0 (221hz), amplitude, consonant

[6]Backness and rounding harmony pattern together in a variety of languages, including Yawelmani (Cole and Kisseberth, 1995b; Krämer, 2003, and others).

duration (60ms), and vowel duration (250ms), then spliced together to form the nonce words.

A list of the vowel combinations used to form the nonce word stimuli is given in Table 4.2. Each nonce word contained the vowel [a]; in the local conditions, the non-[a] vowels were adjacent to each other, and in the non-local conditions, the non-[a] vowels were non-adjacent. Additionally, the non-[a] vowels either agreed or disagreed with respect to backness/rounding and either agreed or disagreed with respect to height. The order in which the non-[a] vowels appeared was counterbalanced. This resulted in a total of 24 possible vowel sequences. These were combined with the consonants [h], [k], and [g] to form the nonce-word stimuli; each nonce word contained at most one of each of the consonants, in each possible order, resulting in six items for each vowel sequence, and 144 items total.

4.4.1.2 Subjects and Task

Subjects were 36 native speakers of North American English, students in introductory linguistics courses recruited by offering extra credit points. Data from three subjects was thrown out because they failed to respond on more than 10% of the trails (responses were cut off after 1500ms), so data from 33 subjects was used.

Subjects were presented with a nonce-word stimulus, followed by either [i] or [u], with an ISI of 750ms. They were asked to judge whether or not they had heard the target vowel in the nonce word. Because the target vowels were reliably either [i] or [u], the task is designed to probe recognition and recall of backness/rounding. Stimuli were presented, randomized, in blocks, with a total of 288 trials per block. There were two blocks, with short within-block breaks every 72 trials. Subjects heard each item a total of four times, twice with [i] as the target and twice with [u] as the target.

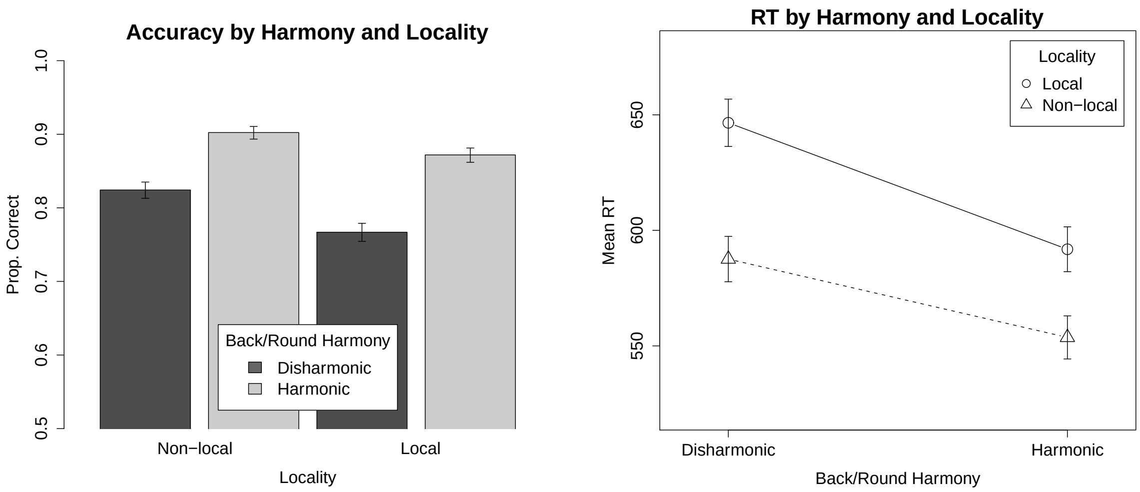

Figure 4.5. Accuracy and response time as a function of back/round harmony and locality. Error bars represent 95% confidence intervals.

4.4.2 Results

As seen in Figure 4.5, subjects were faster and more accurate to identify that they had heard the target vowel when the non-[a] vowels in the nonce word agreed for backness/rounding, regardless of whether those vowels were adjacent or not.

A Linear Mixed Effects Model with log-transformed response time as the dependent variable and random effects for both subject and item found that, for conditions where the non-[a] vowels were adjacent, the effect of backness/rounding harmony was significant ($t = 4.32$, $p < 0.001$).[7] The effect of backness/rounding harmony was also significant for conditions where the non-[a] vowels were non-adjacent ($t = –2.48$, $p < 0.05$).[8] There was no significant interaction between locality and harmony ($t = –1.3$, $p > 0.05$).

A Mixed Logit Model with correct responses as the dependent variable and random effects for both subject and item found similar results for accuracy. For conditions

[7]Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

[8]Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

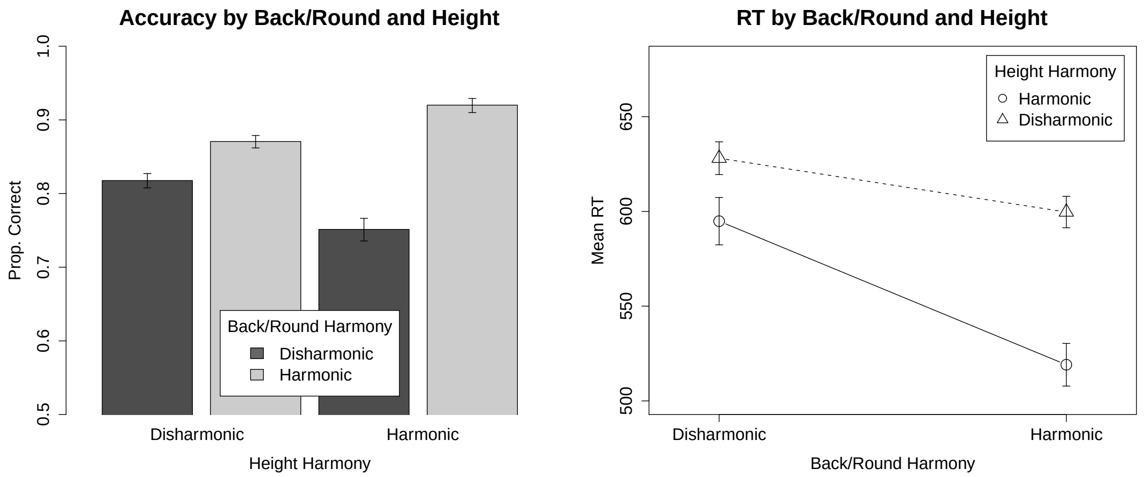

Figure 4.6. Accuracy and response time as a function of back/round harmony and height harmony. Error bars represent 95% confidence intervals.

where the non-[a] vowels were adjacent, the backness/rounding harmony had a significant effect ($z = -7.614$, $p < 0.001$).[9] The effect of harmony was also significant in conditions where the non-[a] vowels were non-adjacent ($t = 6.806$, $p < 0.001$).[10] Again, the interaction between locality and harmony was not significant ($z = 0.368$, $p > 0.05$).

There was also an interaction between back/round harmony and height harmony, echoing the results of the previous experiment. As seen in Figure 4.6, subjects were faster and more accurate to identify the target vowel in back/round harmonic words than in back/round disharmonic words, regardless of height harmony. However, this effect was more dramatic in height-harmonic words than in height-disharmonic words.

A Linear Mixed Effects Model with log-transformed reaction time as the dependent variable and random effects for subject and item found that subjects were significantly faster on back/round harmonic words than back/round disharmonic words in both

[9]Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

[10]Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

height harmonic contexts ($t = 8.33$, $p < 0.001$)[11] and height disharmonic contexts ($t = -2.86$, $p < 0.01$)[12], with a significant interaction between back/round harmony and height harmony ($t = -2.73$, $p < 0.001$).

A Mixed Logit Model with correct responses as the dependent variable and random effects for subject and item found that subjects were significantly more accurate on back/round harmonic words than back/round disharmonic words in both height harmonic contexts ($z = -11.499$, $p < 0.001$)[13] and height disharmonic contexts ($z = 5.095$, $p < 0.001$)[14], with a significant interaction between back/round harmony and height harmony ($z = 6.621$, $p < 0.0001$).

4.4.3 Discussion

The results of this experiment support the hypothesis that harmony is advantageous because it boosts the perceptual salience of feature contrasts. Subjects were consistently faster and more accurate at recognizing the target vowel when the nonce word contained another vowel with the same backness/rounding specification. Furthermore, these results support the hypothesis that this advantage obtains even among non-adjacent vowels.

The advantage of back/round harmony in height harmonic contexts might be expected simply as an effect of the task, because it results in total identity between the non-[a] vowels. A nonce word with two instances of [i] or two instances of [u] might increase the probability of a correct identification for reasons independent of

---

[11] Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

[12] Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

[13] Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

[14] Significance threshold reflects p-value adjusted for multiple comparison using a Bonferroni correction.

harmony. In nonce words like [hikiga], a correct identification of *either* non-[a] vowel is sufficient to successfully perform the task — accurately identifying the initial [i] can result in a correct response, and accurately identifying the medial [i] can also result in a correct response. Improved performance on these conditions does not necessarily indicate that harmony in general is beneficial.

This is not true, however, of height disharmonic conditions. In a nonce word like [hekiga], correct identification of the initial [e] does not independently bear on the question of whether [i] was contained in the word, and cannot in and of itself result in a correct response. The results in Figure 4.6 show that even in these contexts, subjects were significantly faster and more accurate at recognizing the target [i] or [u]. This suggests that the presence of another vowel that agrees with the target with respect to backness/rounding (the features that distinguish the two possible targets from each other) aids in the correct perception of the target.

The results shown in Figure 4.5 demonstrate that this effect is not restricted to adjacent vowels. Subjects were faster and more accurate in responding to back/round harmonic words than disharmonic words, regardless of whether the non-[a] vowels were adjacent or non-adjacent. There was no significant interaction of back/round harmony and locality, indicating that the advantage of harmony over disharmony obtains equally among local and non-local segments.

Words where the non-[a] vowels were non-adjacent were faster and more accurate across the board, likely because in those cases the most recent informative vowel is temporally closer to the subsequent presentation of the target vowel. In a trial like [hukago…u], for example, less time passes between the potentially-informative [o] and the presentation of the [u] than in a trial like [hukoga…u]. Crucially, this effect applied to both harmonic and disharmonic words. The lack of interaction indicates that any advantage that recency provides does not impact the advantage of harmony.

4.4.4 Summary

This section presented the results of a phoneme recall study designed to test the hypothesis that harmony is perceptually advantageous because it increases the salience of perceptually weak contrasts. Furthermore, the experiment tested the hypothesis that this advantage also applies to harmony among non-adjacent segments as well as harmony among adjacent segments

Subjects were faster and more accurate to recognize the target vowel in nonce words which were harmonic for backness/rounding than in disharmonic words. This was equally true both for words where the harmonic vowels were adjacent and for words where the harmonic vowels were non-adjacent, supporting the claim that the perceptual advantages of harmony apply to transparent harmony as well.

4.5 Conclusion

In this chapter, I have presented experimental evidence that harmony is perceptually advantageous — in both discrimination and phoneme recall tasks, subjects performed better on harmonic nonce-word stimuli than on disharmonic stimuli. Crucially, the experiment presented in Section 4.4 demonstrates that this is true not only for local harmony, but also for harmony among non-adjacent segments.

The results presented here contradict the common assumption that a theory of phonology that adheres to strict locality is necessarily a more phonetically grounded theory. If harmony is perceptually (as well as articulatorily) motivated, and if that perceptual motivation obtains even among non-adjacent elements, there is no sense in which strict locality provides a greater degree of phonetic grounding. A theory of Gen which permits gapped representations (and a theory of Con which permits constraints that may prefer them) need not sacrifice grounding in order to do so.

This chapter has not, however, addressed the question of whether strict locality is *phonologically* viable or desirable. For further discussion of strict locality in phonology, and a comparison between the proposal advanced in this dissertation and accounts of transparency which maintain strict locality, see Chapter 6.

CHAPTER 5

VARIATION AND TRIGGER COMPETITION IN FINNISH

5.1 Introduction

In languages like Finnish and Hungarian, lower non-participating vowels are more likely to be opaque to harmony along the front/back dimension than higher non-participating vowels. There are several potential sources of explanation for this asymmetry. Nevins (2004) and Kaun (1995) suggest that the higher sonority (and hence longer duration) of lower vowels means that they constitute a greater interruption to the harmony domain. Benus (2005); Benus and Gafos (2007) propose that the quantal properties of vowels are responsible — high vowels' greater tolerance of transparency is attributable to their greater ability to undergo coarticulation without affecting their categorical perceptual identity.

In previous chapters, I have presented a theory of transparency and opacity in harmony which views the choice between competing harmony triggers — local on the one hand, and dominant on the other. Under this approach, the asymmetries seen in Finnish and Hungarian receive their explanation from lower vowels' status as preferential triggers of front/back harmony; low vowels are more likely to be opaque because they are under more pressure to instigate their own feature domain.

Any of these explanations is sufficient to explain the fact that, in Finnish and Hungarian, higher vowels are more likely to be transparent than lower vowels. However, these three approaches make divergent predictions about the role of the phonological context beyond the transparent/opaque vowel itself. In particular, in a V_1–V_2–V_3

sequence, where V_2 is a potentially transparent or opaque non-participant, do properties of V_1 influence the likelihood that V_2 will be transparent?

A sonority-based explanation for the behavior of V_2 predicts that V_1 should not have any influence over the outcome, since V_1 has no impact on the sonority of V_2. An explanation based on the quantal properties of V_2 predicts that a more articulatorily extreme V_1 (in the case of front/back harmony, higher) should increase the likelihood of transparency — a more extreme V_1 will have a more pronounced coarticulatory influence on V_2, and a V_2 that has undergone more coarticulation for the harmonizing feature value is more likely to be transparent.

An explanation based on trigger competition predicts that, just as opacity is more likely when V_2 is a preferred trigger than when it's not, transparency is more likely when V_1 is a preferred trigger than when it's not. Because weakly cued segments are preferred triggers, this means that a weakly cued V_1 — in the case of front/back harmony, lower — should increase the likelihood of transparency. This is precisely the opposite of what an explanation based in the quantal properties of transparent and opaque vowels predicts.

The data to distinguish between these approaches is not available if we are only concerned with categorical phonological phenomena. However, considering gradient data, as in cases of phonological variation, can provide some insight. In particular, variable transparency and opacity in Finnish loanwords can provide a fertile testing ground for determining whether (and how) the properties of V_1 influence the likelihood of transparency. Because V_2 in Finnish loanwords can be either transparent or opaque, it's possible to observe the factors that increase the probability of one outcome over the other.

In this chapter, I present the results of a nonce-word experiment showing that, in Finnish, the predictions of trigger competition are borne out — transparency is more likely *from* a perceptually impoverished trigger, and less likely *across* one. Section

5.2 provides background information on variable transparency and opacity in Finnish loanwords. Section 5.3 presents the experiment, and Section 5.4 models those results.

5.2 Variable Transparency in Finnish

As discussed in Chapter 2, Finnish has a well-described system of harmony along the front-back dimension (Kiparsky, 1973, 1981; Goldsmith, 1985; Ringen, 1988; Steriade, 1987; Vago, 1988, and others). The inventory is as in (139) — low vowels and round vowels are paired along the front/back dimension, but front unrounded vowels are unpaired (they lack back counterparts [ɯ,ɤ]).

(139) **Finnish Vowel Inventory**

i	y		u
e	ø		o
æ		a	

Within native roots, paired front and back vowels may not co-occur. Furthermore, suffix vowels alternate to take on the feature specification of the root (140).

(140) **Finnish Back Harmony**

a. pøytæ-næ ‘table-ESS’

b. pouta-na ‘fine weather-ESS’

The unpaired front vowels [i,e] have often been described as neutral, since they may co-occur with both front and back vowels. When a neutral vowel intervenes between two paired vowels, it behaves as transparent — both root co-occurrence restrictions and suffix alternations continue to apply to the flanking paired vowels, regardless of the neutral intervenor. The suffix vowels in (141) all must agree with the non-neutral root vowels — for example, a back initial vowel in (141a) requires a back suffix, despite the intervening front [i].

(141) **Finnish Transparent [i,e]**

a.	tunte-vat	'feel-3PL'
b.	puhe-han	'speech-EMPH'
c.	tsaari-na	'car-ESS'
d.	ukit-han	'grandfathers-EMPH'
e.	palttina-lla-ni-han	'with my linen cloth, as you know'
f.	værttinæ-llæ-ni-hæn	'with my spinning wheel, as you know'
g.	luo-da-kse-ni-ko	'for me to create?'
h.	lyø-dæ-kse-ni-kø	'for me to hit?'

In the native vocabulary, neutral vowels are the only source of disharmony. However, loanwords are not forced to harmonize when they are adapted — as a result, front and back paired vowels can co-occur in borrowed roots (142). For example, in (142a), an initial back [u] is followed by a front [æ]; in (142b), an initial front [y] is followed by a back [a], and so forth.

(142) **Finnish Disharmonic Loanwords**

a.	vulgæːri	'vulgar'
b.	tyranni	'tyrant'
c.	afæːri	'affair'
d.	analyːsi	'analysis'

While vowels in loanwords fail to undergo harmony, they still participate as triggers in harmonic suffix alternations (143). Loans where all stem vowels are back uniformly take back suffixes(143a–d), as do disharmonic stems consisting of a front–back sequence (143e–f). Loans consisting only of front vowels consistently take front suffixes (Ringen and Heinämäki, 1999).

(143) **Suffix Harmony with Loanwords**

a.	tabasko-a	'Tabasco'	(*tabasko-æ)
b.	jurtt-a	'yurt'	(*jurtt-æ)
c.	rotunda-sta	'rotunda'	(*rotunda-stæ)
d.	saluuna-sta	'saloon'	(*saluuna-stæ)
e.	syntaksi-a	'syntax'	(*syntaksi-æ)
f.	symptomi-a	'symptom'	(*symptomi-æ)

In loanwords with back–front sequences, however, the value of the suffix vowel is variable; it may either be front or back (144). For example, the partitive singular suffix may surface as either back [-a] or front [-æ] with a back–front word like [afæːri].[1] In other words, the non-neutral front vowels in these loanwords may either be transparent (surfacing with a back suffix) or opaque (surfacing with a front suffix).

(144) **Variability in Disharmonic Loanwords**

a.	hieroglyfi-a	∼	hieroglyfi-æ	'hieroglyph'
b.	analyːsi-a	∼	analyːsi-æ	'analysis'
c.	marttyːri-a	∼	marttyːri-æ	'martyr'
d.	sutenøːri-a	∼	sutenøːri-æ	'pimp'
e.	jongløːri-a	∼	jongløːri-æ	'juggler'
f.	amatøːri-a	∼	amatøːri-æ	'amateur'
g.	miljonæːri-a	∼	mijonæːri-æ	'millionaire'
h.	afæːri-a	∼	afæːri-æ	'affair'

Furthermore, the rate of selecting different variants is not constant across all phonological conditions. Ringen and Heinämäki (1999) conducted a study of 50 native speakers of Finnish, eliciting suffixed forms for disharmonic loanwords. The same speakers were tested twice, with one month between sessions. The table in

[1]A final epenthetic [i] is found when consonant-final loanwords are adapted; in loanwords, [i] is transparent, as it is in the native vocabulary.

(145) shows, for each item, the proportion of subjects whose choice of suffix — front or back — was inconsistent (either within a single session or between the two sessions). Because the vast majority of stable responses were front suffixes (opacity), an increase in variability means an increase in the likelihood of transparency.

(145) **Proportion of subjects who showed variability**

High	($\bar{x}$ =*.35*)
hieroglyfi	.26
analy:si	.48
martty:ri	.32
Mid	($\bar{x}$ =*.18*)
sutenø:ri	.10
jonglø:ri	.18
amatø:ri	.26
Low	($\bar{x}$ =*.12*)
miljonæ:ri	.04
hydrosfæ:ri	.08
afæ:ri	.25

Ringen and Heinämäki did not provide any statistical analysis, but a logistic regression on their data (as given above) shows that the effect of height is significant (Estimate=0.7091, SE=1.512, z=4.689, p=2.74e-06). The higher the vowel, the higher the likelihood that subjects would entertain the possibility of transparent harmony (selection of a back suffix) at least some of the time.

5.3 The Experiment

Of present interest are precisely the type of data that Ringen and Heinämäki (1999)'s experiment is concerned with — disharmonic loans with back–front vowel sequences. The generalization that higher vowels are more likely to be transparent than lower vowels is, as discussed in the introduction to this chapter, equally well explained by a number of different theories of transparency and opacity. Because Ringen and Heinämäki used real loanwords as their stimuli, their results are of limited use in determining what other factors influence the choice between transparency and opacity — the height and number of triggering back vowels, for instance, is not sufficiently manipulated to provide insight into its influence.

The present experiment uses nonce loanwords instead of real loanwords, to permit more detailed manipulation. Subjects are presented with disyllabic back–front loanwords, and asked to make a binary choice between a front suffix (opaque harmony) and a back suffix (transparent harmony). Height and length of both V_1 and V_2 were manipulated, providing the necessary data to test the diverging predictions of various explanations for asymmetries in transparency and opacity.

5.3.1 Hypothesis

If the choice between transparency and opacity is a competition between potential harmony triggers, then the factors that contribute to the strength of *each* of those triggers is predicted to influence the outcome.

In other words, in a V_1–V_2–V_3 sequence, where V_2 is a potentially transparent or opaque non-participant, V_1 and V_2 are competing to determine the feature specification of V_3. In a disharmonic back–front loanword in Finnish, selection of a back suffix represents a victory for V_1 (transparency), while selection of a front suffix represents a victory for V_2 (opacity). If V_2 is a better trigger, then, it should be more likely to triumph than if it is a poor trigger, and hence the probability of opacity should increase. Similarly, if V_1 is a better trigger, it should be more likely to win, and the probability of transparency should increase.

The properties that make a segment a better or worse trigger are related to the factors that motivate harmony in the first place. Since harmony imparts a perceptual advantage by rendering difficult contrasts more salient (as was shown in Chapter 4), the segments that have the most impetus to spread their features are those whose perceptual cues for the relevant contrast are impoverished. Thus, segments with poor cues to the harmonizing feature should be preferred as harmony triggers.

Kaun (1995) shows that this explains asymmetries in trigger strength in various rounding harmony systems; low round vowels are produced with a less prominent rounding gesture than their high counterparts, and these are precisely the vowels which preferentially trigger harmony. The same explanation for triggering asymmetries can also be applied to Finnish — low vowels are less separated along the front/back dimension than their high counterparts, so low vowels should be preferred triggers of harmony along the front/back dimension.

Since the trigger competition model predicts that transparency should be more likely when V_2 is a poor trigger than when it is a preferred trigger, and low vowels are comparatively perceptually impoverished along the front/back dimension (and hence better triggers), **transparency should be more likely when V_2 is high than when it is low**. This explains Ringen and Heinämäki (1999)'s results, which we expect to replicate in the present experiment. Furthermore, the trigger competition model predicts that transparency should be more likely when V_1 is a preferred trigger than when it is a poor trigger, so **transparency should be more likely when V_1 is low than when it is high**.

5.3.2 Methods

5.3.2.1 Stimuli

Stimuli for the experiment consisted of 36 $CV_1C.CV_2$ nonce words; in all conditions, V_1 was back and V_2 was front (and non-neutral). Height and length of both

V_1 and V_2 were manipulated, in a fully crossed design — each possible V_1 (a, a:, o, o:, u, u:) was paired with each possible V_2 (æ, æ:, ø, ø:, y, y:).[2]

The CVC.CV shape was chosen to prevent interactions between initial vowel length and durational effects of stress on non-initial vowels (in CV.CV words only, V_2 shows an increase in duration under V_1 stress; see Suomi et al. 2008 for more details).

The consonants [p,t,k,s,f] were used; the inclusion of [f] in particular was intended to encourage subjects to interpret the nonce items as plausible loanwords, since [f] is not present in the native inventory but is permitted in loanwords. Additionally, a randomly selected half of the medial CC sequences were geminates, and the remainder were heterorganic clusters.

In addition to the test items, 6 additional nonce words were included to serve as catch trials. These again were CVC.CV in shape, using the same consonant inventory, but half consisted of back–back sequences and half consisted of front–front sequences. Because behavior on items of this type is categorical and fully predictable, these items can serve as controls and can be used to determine whether or not subjects are participating in good faith in the task.

5.3.2.2 Task

The experiment was delivered over the internet, using the online survey administration software LimeSurvey. Because the phoneme-to-grapheme relationship in Finnish faithfully represents the relevant vowel contrasts, stimuli were presented orthographically rather than auditorily.

Subjects were presented with a frame sentence with a blank, and asked to make a binary choice between a nonce item with a front suffix and that same item with a back suffix. An example of the task is given in (146). The suffixes used were [-ssa/-

[2] One item (Co:C.Cø:) suffered from a typo (presented as Cø:C.Cø:) — it was analyzed as an additional catch trial. The statistical analysis of the results should still be robust in the face of the missing cell.

ssæ] ("in") and [-sta/-stæ] ("of/from") — both the choice of suffix and the order of presentation of the front/back versions was counterbalanced.

(146) Prosessissa tarvittava kemiallinen liuos valmistetaan[3] ______

a. puktyssa

b. puktyssä

The frame sentences chosen centered thematically around chemistry and chemical combinations, to further encourage subjects to treat the nonce items as novel loanwords. Each nonce word was paired consistently with its own frame sentence, and all subjects saw each item once; this resulted in a survey that took subjects approximately 10 minutes to complete.

To mitigate any prescriptive influence, subjects were explicitly instructed to make their choice based on what sounded best to them, rather than on what they might think is the "correct" choice.

5.3.2.3 Subjects

Subjects were recruited via a post in a Finnish Knitters forum on the knit/crochet based social networking site Ravelry. The survey remained active for approximately two days; 285 subjects started the survey, and 209 completed all questions. Respondents were nearly all female, and reported locations from all over Finland.

Of the 209 subjects who responded to all survey items, data was analyzed for only those who reported that they were native speakers of Finnish and responded incorrectly on at most one of the catch trials. A total of 179 subjects met these criteria.

Informal feedback collected from responses to the original forum post indicated that subjects found the task difficult. A number of respondents indicated that they

[3] "The chemical solution needed in the process is made of..."

felt that the frontness or backness of the vowels in the words was relevant, but many had incorrect intuitions about which vowels were front and which were back. None reported adopting any consistent strategy for determining the correct response.

Though the subjects were not compensated for their time, many reported deriving considerable enjoyment from the task. After the survey was deactivated, a number of potential subjects expressed disappointment at having missed their opportunity to participate.

5.3.3 Results

The variable we are interested in is the proportion of subjects who preferred transparent harmony (versus the proportion of subjects who preferred opaque harmony). Because all of the test items consisted of back–front vowel sequences, selection of a *back* suffix is indicative of transparency, while selection of a *front* suffix is indicative of opacity.

The proportion of transparent responses, as a function of V_1 height and V_2 height, is plotted in Figure 5.1. Overall, subjects showed a preference for opacity — front suffixes were substantially more frequent than back suffixes, across all conditions. Furthermore, both V_1 and V_2 height have an impact on the proportion of transparent responses. In the figure on the left, we see that transparency is more likely with a mid V_1 than with a high V_1, and even more likely with a low V_1. In the figure on the right, we see that transparency is more likely with a high V_2 than with a mid or low V_2, but mid and low V_2 pattern similarly. In Figure 5.2, proportion of transparent responses is plotted as a function of V_1 length and V_2 length. For both V_1 and V_2, subjects were more likely to choose transparency with a short vowel than with a long vowel.

A generalzed linear mixed effects model was fitted, using the `lme4` package (Bates and Maechler, 2010) in R (R Development Core Team, 2009), with back responses as

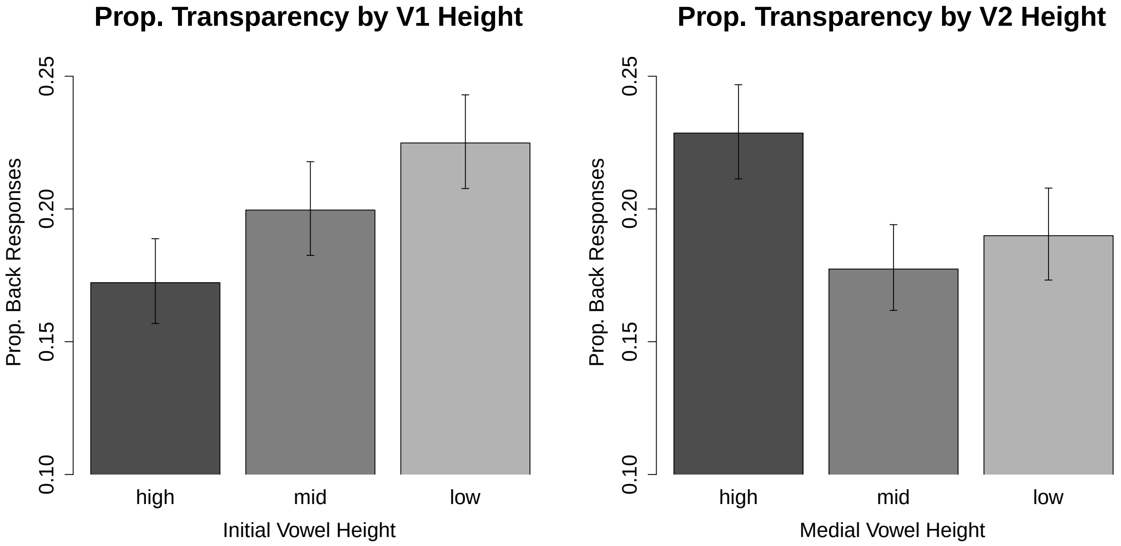

Figure 5.1. Proportion of back suffix responses by initial vowel height (left) and by medial vowel height (right). Error bars represent binomial confidence intervals (95%).

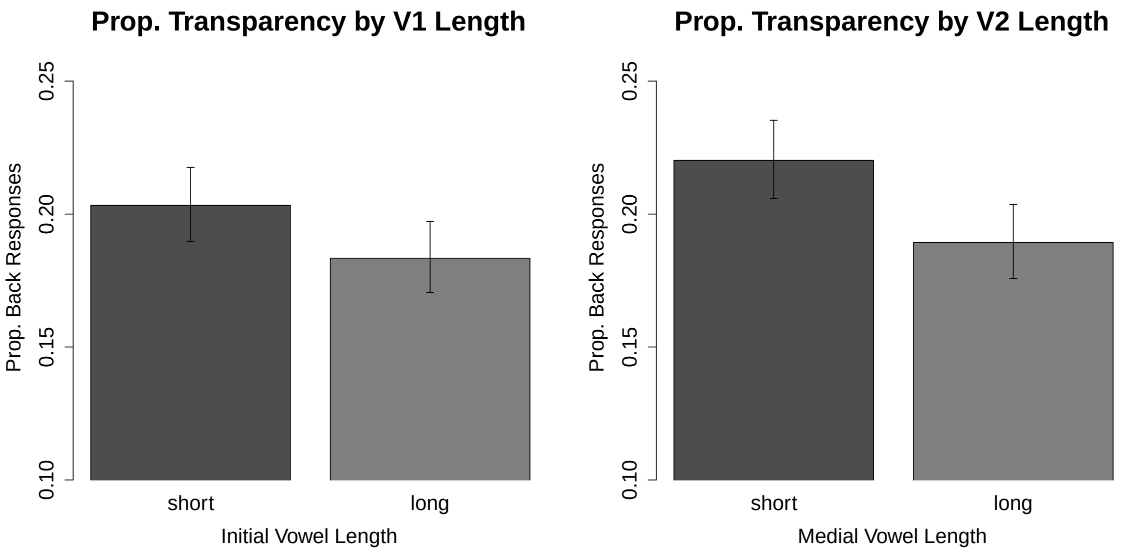

Figure 5.2. Proportion of back suffix responses by initial vowel length (left) and medial vowel length (right). Error bars represent binomial confidence intervals (95%).

the dependent variable, V_1 height, V_1 length, V_2 height, and V_2 length as predictors, and random effects for subjects. Vowel height is treated as an ordered factor — the three levels (low, mid, high) are in an ordinal relationship, as mid vowels are articulatorily and acoustically intermediate between high and low vowels, but there is no a priori reason to believe that the distance between low and mid vowels is the same as the distance between mid and high vowels. As an ordered factor, height is given orthogonal polynomial contrast coding.

The effects of height were significant for both V_1 and V_2; for V_2 height, both the linear ($z = 3.617$, $p < 0.001$) and quadratic ($z = -3.746$, $p < 0.001$) terms were significant; for V_1 height, the linear term was significant ($z = -5.339$, $p < 0.001$), but not the quadratic term ($z = 0.315$, $p = 0.75$). In other words, the likelihood of transparency increases as V_1 height decreases, and this increase is linear. On the other hand, the likelihood of transparency decreases as V_2 height decreases, but this decrease is not linear. The significance of the quadratic component for V_2 height reflects the behavior of mid vowels we saw in Figure 5.1.

Furthermore, the effect of V_2 length was significant ($z = 4.989$, $p < 0.001$), but the effect of V_1 length was not ($z = 1.018$, $p = 0.31$). In other words, the length of V_1 did not have an impact on the likelihood of transparency, but transparency was more likely if V_2 was short than if it was long.

5.3.4 Discussion

The results presented in Section 5.3.3 support the predictions of the trigger competition approach — in particular, the finding that the likelihood of transparency is greater with lower vowels supports trigger competition to the exclusion of other proposed explanations for transparency/opacity asymmetries.

Because lower vowels are more poorly cued along the front/back dimension than their high counterparts, they have a greater impetus to spread their feature value and reap the perceptual rewards of harmony. Because V_1 and V_2 are competing to spread their value for the feature onto V_3, both V_1 and V_2 should be more likely to win when they are preferred triggers than when they are poor triggers. Thus, a victory for V_1 (transparency) should be more likely when V_1 is low than when it is high — similarly, a victory for V_2 (opacity) should be more likely when it is low than when it is high (meaning that, in turn, transparency should be less likely when V_2 is low than when it is high). This is precisely what the experimental results show to be the case in Finnish.

At first glance, the results for vowel length appear to support a sonority-based explanation over one situated in trigger competition — V_1 length had no significant effect, while for V_2, transparency was more likely for short (less sonorous) vowels than for long vowels. If weakly cued segments are better triggers, short vowels should be preferred as triggers; long vowels' increased duration provides more robust cues for vowel features. Therefore, short vowels in V_1 position should result in a greater likelihood of transparency than long vowels, and short vowels in V_2 position should result in a decreased likelihood of transparency.

While the contribution of V_1 length did not reach statistical significance, the numerical trend is in the direction predicted by the competing triggers account — transparency was slighly more likely with a short V_1 than with a long V_1. The results for V_2 length at first appear to be more troubling; the effect is significant, but in the wrong direction. Short vowels should be preferred as triggers, and hence should result in decreased likelihood of transparency, but in fact they are *more* likely to be treated as transparent than their long counterparts. However, these results are expected if we consider the fact that intrinsic segmental properties are not the only source of trigger strength — in fact, crucial to the account of transparency and opacity is the notion that non-local triggers are worse triggers, and increasing degrees of non-locality result in an increasingly bad trigger. Long vowels, which are both durationally and representationally longer than short vowels, result in an increased degree of non-locality. Thus, transparent harmony across a long vowel is *less local* than transparent harmony across a short vowel, and should in fact be less likely, as the results show.

The competing triggers account predicts an effect for V_1 length which does not reach significance in the experimental results. The fact that there is a numerical trend in the predicted direction can't be said to count in *favor* of this approach, but also cannot necessarily be said to count *against* it. On the other hand, the sonority-based approach predicts that no properties of V_1 should have an influence, yet V_1 height is a significant predictor.

A comparison was made between the generalized linear mixed effects model described above (Model 1: V_1 height, V_1 length, V_2 height, and V_2 length as predictors) with one which omitted the V_1 parameters (Model 2: only V_2 height and V_2 length as predictors). Model 1 performed better than Model 2 on AIC (Model 1 = 4556.4, Model 2 = 4579.8), BIC (Model 1 = 4610.3, Model 2 = 4613.5) and log likelihood (Model 1 = -2284.9, Model 2 = -2270.2). A χ^2 test found the difference between the two models to be significant ($\chi^2 = 29.45$, $p < 0.001$).

These model comparison tests penalize including additional factors, but reward the increase in explanatory power that comes from inclusion of significant predictors. Model 1 represents the predictions of a trigger competition account; V_1 length and height are both included, even though V_1 length is not significant, because the theory predicts that they should have an influence over the outcome. Model 2 represents the predictions of a sonority-based account; no V_1 parameters are included, because these are not predicted to have an effect. In comparing these two models, we see that the explanatory value of including V_1 height as a predictor appears to be worth the cost of including V_1 length as a non-significant predictor. In other words, while neither set of predictions is perfectly matched, a trigger competition model provides a better account of the data than a sonority-based account.

The account that Benus (2005); Benus and Gafos (2007) propose, however, does have the ability to predict that the properties of V_1 can influence the choice between transparency and opacity — however, those precitions would be that the effect should be in the opposite direction. Transparency is predicted to be more likely when V_1 is high than when it is low, contrary to the effect found in the experimental results above.

For Benus and Gafos, transparency is in fact a form of participation. A non-participating V_2 is unable to undergo categorical alternation, but does undergo sub-phonemic coarticulation — that is, because an [i] following a back vowel coarticulates with its neighbor, it is less front that an [i] followed by a front vowel. This retraction, in turn, predicts whether V_3 will be front or back; a sufficient degree of sub-phonemic retraction on V_2 induces categorical backness on V_3. Factors which increase the degree to which V_2 coarticulates, then increase the probability of transparency.

Under this account, the asymmetries in V_2 height are due to non-linearities in the relationship between articulation and perception. According to quantal theory, [i] is in a zone of relative perceptual stability as compared to lower front vowels, and can therefore tolerate a greater degree of articulatory variability without affecting its perceptual identity as a front vowel. Lower vowels, because they are closer to a region of comparative instability, can tolerate relatively less articulatory variability if they are to remain perceptually identifiable as front vowels. As a result, a high non-participating V_2 can undergo more articulatory retraction following a back vowel than a mid or low V_2 — because more retraction means a greater probility of transparency, this explains why high vowels are more likely to be transparent. Benus and Gafos use this difference in degree of retraction to explain why hungarian [e] is less likely to be transparent than [i], and also to explain why a sequence of two [i] syllables is optionally opaque (the coarticulatory effect has a diminished influence on the second [i] in the sequence, resulting in a lower likelihood of transmitting backness to a subsequent suffix).

Benus and Gafos do not discuss the effects of V_2 length, but those can also be explained via influence on the degree of coarticulation. Because greater vowel duration means more time to fully reach an articulatory target, a long [i] should undergo less coarticulation, and should be less retracted than a short [i]. In turn, because long vowels will be less retracted, they should be less likely to induce backness on a following vowel and therefore should be less likely to be transparent than their short counterparts — the effect here is similar to distance effect shown by sequences of multiple transparent vowels. This is precisely the effect we see in the Finnish data.

Benus and Gafos also do not discuss the role of V_1, but their explanation for the other effects does provide a way of predicting what the effects should be. Higher vowels are more extreme in their articulation along the front/back dimension than lower vowels — high vowels, therefore, should exert more coarticulatory influence over the following vowel than their low counterparts, since reaching a front target requires greater articulatory movement when coming from a more extreme back vowel. This increased degree of coarticulation, in turn, should mean a greater likelihood of transparency. In other words, transparency should be more likely when V_1 is high than when it is low — in fact, however, Finnish subjects were more likely to choose transparency when V_1 was low than when it was high.

A similar reasoning applies when predicting the effect of V_1 length. Because a longer vowel means that a more extreme articulatory target is possible, a longer V_1 should induce more coarticulation on V_2 than a shorter vowel. This in turn should induce a greater likelihood of transparency, predicting that transparency should be more likely when V_1 is long than when it is short. There was no significant difference between long and short V_1 in the experimental results, however the numerical trend was in the opposite direction from these predictions — short vowels were slightly more likely to induce transparency than long vowels.

A side-by-side comparison of the predictions of the three models with the experimental results is given in (147). All three models successfully account for the V_2 results for height and length, but diverge with respect to V_1. Because V_1 length did not reach significance, it is not particularly useful in choosing among the theories; the lack of effect is predicted by the sonority-based account, and the numerical trend in the data is predicted by the trigger comparison account. V_1 height showed a clear, significant effect, so it is perhaps more useful — and clearly favors a trigger competition approach.

(147) **Attested vs. Predicted Likelihood of Transparency**

	V_1 Height	V_1 Length	V_2 Height	V_2 Length
a. Exp. Results	**hi < lo**	**short = long**	**hi > lo**	**short > long**
b. Sonority	hi = lo ✗	short = long ✓	hi > lo ✓	short > long ✓
c. Quantal Theory	hi > lo ✗	short < long ✗	hi > lo ✓	short > long ✓
d. Trigger Comp.	hi < lo ✓	short > long **?**	hi > lo ✓	short > long ✓

Of the three approaches to explaining asymmetries in transparency and opacity, the trigger competition approach advocated in this dissertation best accounts for the qualitative patterns in the results of this experiment. In particular, this approach uniquely predicts that lower vowels in V_1 should result in a greater likelihood of transparency, an effect which is present and robustly statistically significant in the experimental results.

5.4 Modeling Variable Data

In addition to providing a qualitative explanation of the patterns found in the data, the theoretical architecture presented in Chapter 3 can be used to successfully model the experimental results presented above. The theoretical model presented thus far has assumed *categorical* data — that is, for each input, there is only one possible outcome. The present task is to model variable data; there are multiple attested outcomes for each input, and we are interested in the relative probabilities of each of those outcomes.

Before proceeding, it is important to note that the present discussion follows much of the literature on variation in phonological theory in conflating *inter*– and *intra*–speaker variation. This is primarily a function of the data available — it is far easier to observe data between speakers than to obtain the data necessary for careful study of variation within a single speaker. By conflating the two, the assumption is made that there are no significant sociolinguistic or regional factors driving the inter-speaker variation, and that the probability distribution across a population of speakers is representative of the probability distribution that characterizes individual speakers' behavior.

For the present case, this assumption appears to be broadly supported. Ringen and Heinämäki (1999)'s results were based in part on rates of intra-speaker variation, since they elicited judgments from the same subjects across multiple sessions. The data from this study replicates their result (V_2 is more likely to be treated as transparent when it is high than when it is low), establishing a basic similarity between the inter-subject and intra-subject data. Furthermore, (Ringen and Heinämäki, 1999) found no consistent sociolinguistic source of variability for this phenomenon. These facts conspire to suggest that the conflation of inter– and intra–speaker variation, while less than ideal, is not grossly problematic for the data currently under consideration.

In order to model this variability, it is necessary to add a stochastic component to the grammar. This can be done by perturbing the constraint weights themselves with (normally distributed) noise at evaluation, as in Noisy Harmonic Grammar (Noisy HG); this kind of evaluation noise was first proposed in Stochastic OT by Boersma (1998a); Boersma and Hayes (2001) and implemented in HG in Praat Boersma and Weenink (2008). Alternatively, stochasticity can be accomplished by defining a probability distribution across candidates based on the exponents of their harmony scores, as in Maximum Entropy grammar (MaxEnt; Goldwater and Johnson 2003; Wilson 2006; Jäger and Rosenbach 2006, and others). See e.g. Coetzee and Pater (2008) for some comparison of Noisy HG and MaxEnt; the following discussion uses a MaxEnt model, but nothing crucial hinges on that choice, and Noisy HG should be equally applicable.

To arrive at an optimal set of weights, the data were subjected to a batch learner.[4] The learner receives all the data at once (in this way it differs from online learning algorithms like the HG-GLA; see e.g. Boersma and Pater 2008) and minimizes the Kullback-Liebler divergence (Kullback and Leibler, 1951) between the observed probability distribution and the expected probability distribution (as predicted by MaxEnt) by adjusting the weights of the constraints. To include examples of local harmony, the catch trials were included; they behave more or less categorically, in that the proportions for front and back responses are extremely close to either 0 or 1.

The constraints used, SPREAD(+BACK) and SPREAD(–BACK), are given in (148), with scaling factors for perceptual impoverishment and non-locality given. While the precise numerical values of the constants applied in the scaling factors were unimportant for the categorical data in Chapter 3, here it becomes somewhat crucial — the

[4]The batch learner is scripted by Robert Staubs, and implemented in R.

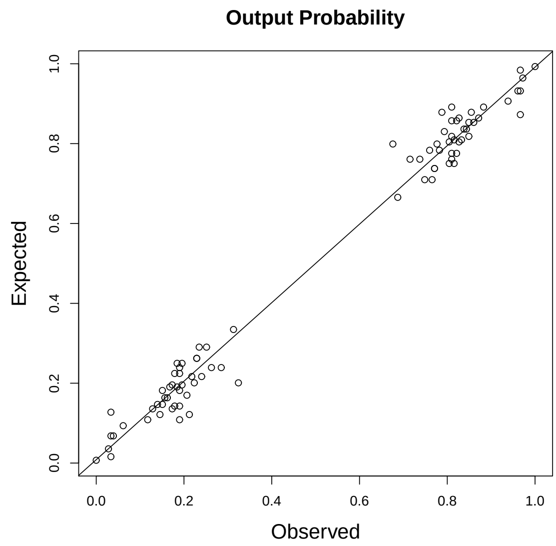

Figure 5.3. Observed probabilities versus probabilities for both back and front outputs, for each of the experimental stimuli. Black line is the slope of a linear regression (R^2 = .9842).

greater the scaling factor, the greater the predicted effect. In principle, these could be learned alongside the constraint weights; for the time being, the numerical values are based loosely on the estimates from the Generalized Linear Mixed Effects model presented in Section 5.3.3.

(148) a. SPREAD(+BACK): Assign +1 for each segment linked to [+BACK] as a dependent.

- **Scaling factor: non-locality**

  For a trigger α and a target β, multiply the reward earned for the dependent segment β by a constant k (such that $1 > k > 0$) for each unit of distance d intervening between α and β.

  – If d = V, $k = 0.25$

  – If d = V:, $k = 0.155$

- **Scaling factor: trigger strength**

For a trigger α, a target β, and a feature F, multiply the reward earned for the dependent segment β by a constant x (such that $x > 1$) for each degree i to which α is perceptually impoverished with respect to ±F.

- If $i = 1$ (mid), $k = 1.25$
- If $i = 2$ (low), $k = 1.5$

b. SPREAD(−BACK): Assign +1 for each segment linked to [−BACK] as a dependent.

- **Scaling factor: non-locality**

  For a trigger α and a target β, multiply the reward earned for the dependent segment β by a constant k (such that $1 > k > 0$) for each unit of distance d intervening between α and β.

  - If $d =$ V, $k = 0.25$
  - If $d =$ V:, $k = 0.155$

- **Scaling factor: trigger strength**

  For a trigger α, a target β, and a feature F, multiply the reward earned for the dependent segment β by a constant x (such that $x > 1$) for each degree i to which α is perceptually impoverished with respect to ±F.

  - If $i = 1$ (mid), $k = 1.18$
  - If $i = 2$ (low), $k = 1.36$

The overall model of the results fits the attested data quite well; Figure 5.3 plots the expected probabilities produced by the MaxEnt model against the observed probabilities from the experimental data. There is a high degree of correlation between the observed and expected results ($R^2 = 0.98$), indicating that the learner is able to use the scalar constraints in (148) to produce a distribution appropriate to the qualitative patterns in the data.

The analysis that follows is concerned only with the choice between transparent and opaque harmony. To achieve non-participation, a faithfulness constraint specific to loanword stems must have a weight sufficiently greater than both the spreading constraints, ensuring that full participation is not an option. Furthermore, ordinary faithfulness constraints are weighted sufficiently lower than both the harmony constraints, ensuring that the suffix vowel will always harmonize with one of the two root vowels. These constraints are not included in the analysis, but it should be noted that in MaxEnt models, some probability is assigned to *every* candidate (even harmonically bounded candidates). In this way, MaxEnt differs from Noisy HG. If included, candidates with full participation or suffix faithfulness would receive a vanishingly small share of the probability distribution; the specific proportions assigned to the viable candidates would therefore be slightly reduced, but this does not affect the analysis in any substantial way.

In the tableaux that follow, constraints and their weights are given to the left of the jagged line. To the right of the jagged line, the harmony scores ($\mathcal{H}$), observed proportion of responses from the experimental data (Obs.), and predicted probability from the MaxEnt model (MaxEnt) are given.

The tableaux in (149) show the analysis of V_1 height effects — V_2 is kept constant. In each of (149a-c), candidate (a) is the form with a front suffix (–ssä/–stä); this earns a reward of +1 on Spread(F). Because this is local harmony, the scaling factor for non-locality is not applied; similarly, because the front vowel is high (and therefore not a preferred trigger), the scaling factor for preferential triggering is not applied, so the reward remains as it is. Candidate (a)'s harmony score of 1.924 is the result of multiplying that reward by the weight of Spread(F) — because it earns no rewards on Spread(B), summing across all constraints does not affect the harmony score.

The (b) candidates are of present interest; these are the candidates where a back suffix has been chosen, and represent transparent harmony. In (149a), candidate (b) earns a reward on SPREAD(B); because it is non-local, the scaling factor for non-locality is applied, and the reward is multiplied by 0.25; because V_1 is high, and therefore not a preferred trigger, the scaling factor for preferential triggering is not applied. In (149b), V_1 is mid, so in addition to being scaled for non-locality, candidate (b)'s reward is multiplied by 1.25 because it is a somewhat preferred trigger. Finally, in (149c), V_1 is low — because it is a preferred trigger, its reward is multiplied by 1.5 (rather than the 1.25 we saw for the initial mid vowel).

The relationship between the three items in (149) can be seen by comparing the harmony scores and observed and predicted probabilities for each of the (b) candidates — transparency (selection of a back suffix) is more harmonic and more likely in (149b) than in (149a), and more likely in (149c) than in (149b); we have successfully modeled the generalization that transparency increases in likelihood as V_1 decreases in height.

(149) a.

pukty	SPREAD(B) 3.297	SPREAD(F) 1.924	$\mathcal{H}$	Obs.	MaxEnt
a. pukty-ssä		+1	1.924	0.804	0.750
b. pukty-ssa	+0.25 (+1 * .25)		0.824	0.196	0.250

b.

kofty	SPREAD(B) 3.297	SPREAD(F) 1.924	$\mathcal{H}$	Obs.	MaxEnt
a. kofty-stä		+1	1.924	0.749	0.710
b. kofty-sta	+0.313 (+1 * 1.25 * .25)		1.030	0.251	0.290

c.

sakfy	SPREAD(B) 3.297	SPREAD(F) 1.924	$\mathcal{H}$	Obs.	MaxEnt
a. sakfy-stä		+1	1.924	0.687	0.665
b. sakfy-sta	+0.375 (+1 * 1.5 * .25)		1.236	0.313	0.335

The tableaux in (150) show the effect of V_2 height — here, V_1 is held constant. The (b) candidates, representing transparent harmony, all receive a reward on SPREAD(B) that is scaled for both non-locality (multiplied by 0.25) and its status as a (low) preferential trigger (multiplied by 1.5).

It is the (a) candidates which are now of present interest. In (150a), as in the examples above in (149), V_2 is high, and so its reward on SPREAD(F) is not modified by any scaling factor. In (150b), V_2 is mid, and thus a somewhat preferred trigger, so its reward is multiplied by 1.18 in accordance with the scaling factor for preferential triggering. In (150c), V_2 is low, and thus a preferred trigger, so its reward is now multiplied by 1.36; because harmony from V_2 is local, none of the rewards for the (a) candidates in (150) are scaled for non-locality.

Again we can observe the success of the analysis by examining the relationship between the harmony scores and observed and expected probabilities for back suffix selection (transparency; the (b) candidates) across the tableaux in (150). As we expect, transparency is more harmonic and more likely in (150a) than in (150b), and likewise more harmonic and more likely in (150b) than in (150c), correctly modeling the generalization that the likelihood of transparency increases as V_2 height increases.

(150) a.

sakfy	SPREAD(B) 3.297	SPREAD(F) 1.924	$\mathcal{H}$	Obs.	MaxEnt
a. sakfy-stä		+1	1.924	0.687	0.665
b. sakfy-sta	+0.375 (+1 * 1.5 * .25)		1.236	0.313	0.335

b.

fassö	SPREAD(B) 3.297	SPREAD(F) 1.924	$\mathcal{H}$	Obs.	MaxEnt
a. fassö-stä		+1.18 (+1 * 1.18)	2.270	0.771	0.737
b. fassö-sta	+0.375 (+1 * 1.5 * .25)		1.236	0.229	0.262

c.

tafsä	SPREAD(B) 3.297	SPREAD(F) 1.924	$\mathcal{H}$	Obs.	MaxEnt
a. tafsä-ssä		+1.36 (+1 * 1.36)	2.61	0.776	0.799
b. tafsä-ssa	+0.375 (+1 * 1.5 * .25)		1.236	0.223	0.200

The effects of V_2 length are shown in (151), and are the result of scaling for non-locality. Here V_1 is held constant, as is V_2 height. In both (151a) and (151b), V_2 is high, and so its reward on SPREAD(F) is not scaled for preferential triggering; because harmony for V_2 is local, rewards on SPREAD(F) are also not scaled for non-locality. In both (151a) and (151b), V_1 is low, so its reward is multiplied by 1.5 according to the scaling factor for preferential triggering. The difference between (151a) and (151b) is in the scaling factor for non-locality. In both cases, back suffix selection represents non-local harmony, so the reward for SPREAD(B) is diminished. In (151a), transparent harmony means skipping a short vowel; the reward is mulitiplied by 0.25 because of that non-locality. However, in (151b), transparent harmony means skipping a long vowel — because this is an increased degree of non-locality, the reward is multiplied by 0.155 instead. The results is that transparent harmony across a short vowel is more harmonic and more likely than transparent harmony across a long vowel.

(151) a.

sakfy	SPREAD(B) 3.297	SPREAD(F) 1.924	$\mathcal{H}$	Obs.	MaxEnt
a. sakfy-stä		+1	1.924	0.687	0.665
b. sakfy-sta	+0.375 (+1 * 1.5 * .25)		1.236	0.313	0.335

b.

paftyy	SPREAD(B) 3.297	SPREAD(F) 1.924	$\mathcal{H}$	Obs.	MaxEnt
a. paftyy-stä		+1	1.924	0.715	0.761
b. paftyy-sta	+0.233 (+1 * 1.5 * .155)		0.767	0.285	0.239

Crucially, the model presented in this dissertation is able to successfully model the qualitative patterns found in the Finnish data. Because the effects are fairly small, it would be possible to produce a probability distribution with a similarly high R^2 value using unscaled versions of SPREAD(B) and SPREAD(F) (with the addition of a constraint specifically penalizing non-locality). However, despite producing probabilities that correlate fairly well with the observed data, this constraint set would be unable to model the observed *patterns* — all transparent candidates would receive the same probability, regardless of V_1 height or V_2 height or length. The model presented here doesn't just closely match the attested probabilities, it matches the qualitative patterns that characterize the choice between transparency and opacity in Finnish.

5.5 Conclusion

In this chapter, I have presented the results of a nonce word experiment designed to test the factors influencing the relative frequencies of transparency and opacity in Finnish loanwords. The results replicate the findings of a previous study showing that subjects are more likely to choose the transparent outcome when V_2 is high than when it is low. Furthermore, the present study finds that V_2 length and V_1 height also influence the choice — subjects were more likely to choose transparency when V_2 was long than when it was short, and they were more likely to choose transparency when V_1 was low than when it was high.

The results of this study support the hypothesis that the choice between transparency and opacity is competition between potential harmony triggers — in particular, the finding that a lower V_1 results in a greater likelihood of transparency is consistent with the claim that perceptually impoverished vowels are preferred triggers, and that preferred triggers are more likely to induce transparency. This finding is inconsistent with previous explanations for the attested V_2 effects — deriving those effects from the sonority of V_2 predicts no role for V_1 whatsoever, and deriving those effects from the ability of a lower V_2 to undergo a greater degree of coarticulation predicts that V_1 effects should pattern in the opposite direction.

Furthermore, I have shown that the theoretical architecture presented in Chapter 3 is able to successfully model the attested data, using a Maximum Entropy model to implement variation in Harmonic Grammar. A batch learner was able to arrive at a set of weights for the two harmony constraints that, when modified by the appropriate scaling factors, produces a probability distribution across candidates which matches the observed probabilities both quantitatively (high R^2) and qualitatively (patterns with the attested effects).

CHAPTER 6

COMPARISON WITH ALTERNATIVES

6.1 Introduction

The issue of transparency in vowel harmony has enjoyed a rich history in the literature, dating back to early work on autosegmental phonology (Clements, 1976; Vago, 1976, and others) and continuing through into work in Optimality Theory. The question of transparency — whether it exists, and how best to model it — remains an open one. In this chapter, I address a number of alternative approaches to modeling transparency and opacity, comparing them with the competing triggers approach advanced in this dissertation and ultimately showing that the present approach provides better empirical coverage.

While discussing each extant proposal in detail would prove an excessively lengthy endeavor, I will discuss a number of approaches more generally. I divide previous analyses into two main types: those which attempt to preserve strict locality, addressed in Section 6.2, and those which allow for non-local interactions, addressed in Section 6.3.

In Section 6.2, in addition to addressing the pitfalls of various analyses which maintain strict locality, I address the issue of whether or not such an analysis is desirable. While considerable work has been devoted to providing strictly local analyses of transparent harmony, relatively less attention has been given to the reasons for doing so in the first place. In section 6.2.4, I examine the arguments in favor of strict locality, ultimately concluding that there is no *a priori* reason to prefer a strictly local analysis.

Instead, analyses should be evaluated based on empirical coverage. I show that the competing triggers proposal provides a better account of the available data than the other proposals discussed in this chapter.

6.2 Approaches Maintaining Strict Locality

Strict locality — the requirement that phonological processes (assimilation processes in particular) must involve elements that are *adjacent* at some level of representation — is at odds with the existence of transparent vowels in harmony. Appealing to different levels of representation is of limited use; while a tier-based conception of locality restrictions (elements must be adjacent on their tier) can account for transparent consonants, there is no principled way to divide tiers that places participating vowels on one tier and transparent vowels on another.

A number of approaches to harmony have sought to reconcile the conflict between strict locality and transparency, providing accounts of transparent harmony which strive to maintain only local interactions. In this section, I address three ways of doing so; appealing to derivational opacity (Section 6.2.1), which maintains strict locality in the application of the phonological rule, but not on the surface; interpreting coarticulation as participation (Section 6.2.3), which maintains surface locality, and relying on underspecification of transparent segments (Section 6.2.2), which maintains strict locality in the form of a prohibition on crossed autosegmental association lines. Each of these approaches encounters problems accounting for the empirical facts; in light of these difficulties, I address the overall desirability of strict locality, arguing that there is no *a priori* reason to prefer a strictly local analysis (6.2.4).

6.2.1 Transparency as Derivational Opacity

One of the oldest approaches to the problem of non-local harmony is to treat surface transparency as an instance of derivational opacity — harmony applies strictly locally at one stage of a phonological derivation, with transparent segments as undergoers, but a subsequent phonological process reverts those transparent segments to their original feature specification. This type of analysis can be found in rule-based approaches to harmony (Bach, 1968; Clements, 1976, and others) as well as OT-based analyses (Walker, 1998; Bakovic and Wilson, 2000, and others).[1]

This approach does not maintain strict locality at the surface level, but does maintain it at the level of the phonological process. At the point at which vowel harmony applies, it applies strictly locally — at the first step in (152), the initial back [ɔ] induces backness on *all* subsequent vowels, including [i:]. This results in an intermediate form with strictly local harmony, but an ill-formed segment [ɯ]. Subsequently, in the second step in (152), a general fronting rule (or well-formedness condition) fronts that [ɯ] to [i], undoing the effects of harmony on that segment.

(152) i Harmony: zɔfiːr-bɛn $\longrightarrow$ zɔfɯːr-bɔn

ii Fronting: zɔfɯːr-bɔn $\longrightarrow$ zɔf<u>iː</u>r-bɔn

This results in a surface form that shows a harmonic relationship among non-adjacent elements. Here the ordering of the harmony rule and the general fronting rule is crucial — because the fronting rule follows the harmony rule, it can undo its effects, resulting in counterfeeding opacity. This kind of effect follows straightforwardly from the architecture of rule-based phonology; adopting a similar approach in OT requires recruiting various devices for handling opacity. Walker (1998) accomplishes this by using Sympathy Theory (McCarthy, 2003b), while Bakovic and Wilson (2000) rely instead on Targeted Constraints (Wilson, 2000).

The problem with this approach lies in the motivation for the process that undoes harmony on transparent segments. In Hungarian, the process of fronting [ɯ] to [i]

[1]Implementing a theory like this in OT requires additional machinery to handle derivational opacity; Walker (1998), for example, adopts Sympathy Theory (McCarthy, 2003b).

is independently motivated — it enforces general inventory restrictions prohibiting non-low back unrounded vowels. An analysis of this type successfully accounts for the facts in languages like Hungarian, where a segments' exclusion from harmony correlates with the language-wide absence of its harmony counterpart.

However, transparency of harmonically paired vowels, as in Finnish disharmonic loanwords, does not receive a successful analysis under this approach. At the first step in (153), the initial back [a] initiates (strictly local) backness harmony, making all subsequent vowels back. What is required now to produce the (attested) transparent [y:] is a general rule fronting [u] to [y] — and herein lies the problem. Such a rule is not motivated in Finnish, nor is it motivated by general principles of markedness; the backness and rounding of [u] are features which enhance each other, while the combination of frontness and rounding on [y] are acoustically antagonistic. If anything, the opposite process — backing of [y] — would be motivated.

(153) i Harmony: marttyːri-æ ⟶ marttuːrɯ-a

ii Fronting: marttuːrɯ-a ⟶ marttyːri-a (*)

OT-based implementations of this approach face the same problem as the rule-based implementation — something must motivate the step of fronting [u] to [y], but there is no available markedness constraint to exert this preference. If there is a constraint motivating [u] to [y] fronting, we would expect to see its effects generally — even restricting the domain of discussion to loanwords, we would expect [u] to always surface as [y], whether or not it was created by harmony.

Because the present account derives transparency from explicitly non-local interactions rather than derivational opacity, it does not suffer from this problem. No constraint preferring the fronting of [u] to [y] is required, because [y] is skipped by harmony. As discussed in Chapter 2, its non-participation can be attributed to faithfulness (specific to loanwords, perhaps) rather than markedness. Furthermore, as was shown in Chapter 5, the theoretical architecture presented in this dissertation is able to successfully account for the option of transparency in Finnish loanwords.

6.2.2 Underspecification

Another way of preserving strict locality in the phonological processes driving harmony is to posit a level of representation at which seemingly non-local elements are in fact adjacent. In accounting for transparent vowels, the task becomes determining what level of representation that is — by what set of criteria are the vowels flanking a transparent vowel adjacent to each other?

A number of locality parameters which are successful in other areas — for example, restrictions on skeletal or prosodic adjacency (see e.g. Steriade 1987) — encounter difficulty with transparent vowels; the flanking vowels which interact across a transparent vowel are not adjacent on the vowel tier, nor are they adjacent with respect to any particular prosodic representation. An alternate method of defining adjacency, advanced by Clements (1976); Kiparsky (1981); Archangeli and Pulleyblank (1994, and others), is to make reference to feature-bearing units in determining locality. That is, two segments can be considered adjacent if they each bear some value of the relevant feature, and are not separated by any other segments which also bear that feature.

In this approach, strict locality takes the form of an inviolable prohibition on crossed autosegmental association lines. Segments may be skipped so long as they are do not bear a value for the spreading feature — that is, so long as they are unspecified. This proposal builds on the claim that "only *irrelevant* segments may intervene between focus and determinant in phonological rules" (Jensen, 1974, p. 680). A segment is relevant if it is specified along the feature dimension referenced in the spreading process; it is irrelevant if it is not.

Kiparsky (1981) presents an analysis harmony in Hungarian which makes use of this approach to non-locality. The non-participation of [i] and [e] is derived through a language-wide prohibition on non-low unround back vowels; this also determines the surface outcome of unspecified vowels (they will be filled in by default as front). An [i] or [e] which is underlyingly unspecified along the front/back dimension (I and E, using Kiparsky's notation) will be transparent to backness harmony; skipping it does not violate the line crossing prohibition, because it is unassociated. This is shown in (154a); a default specification rule subsequently specifies the transparent vowel as front.

An underlyingly specified [i] or [e], on the other hand, will be opaque. Because it is associated with its own autosegment, skipping it would violate the line crossing prohibition; the suffix vowel receives default specification instead. This is shown in (154b).

(154) a. **Transparency**

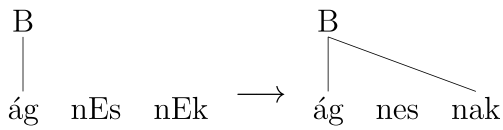

b. **Opacity**

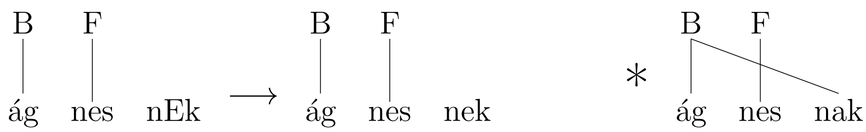

An analysis of transparency as underspecification closely relates the non-participation of transparent segments to the set of contrasts available in the inventory. In Hungarian — and, indeed, in many cases of transparency — the neutral vowels do not need to be underlyingly specified, because they are not contrastive. Well-formedness conditions that apply language-wide in Hungarian mean that the backness specification of an unrounded vowel is predictable; they will be predictably front. It therefore follows that they should be able to be unspecified underlyingly, and should therefore be able to be transparent.

Connecting contrast to transparency in this way makes clear predictions about when vowels can and cannot be treated as transparent — they can be treated as transparent when they do not contrast with respect to the harmonizing feature, but they should never behave as transparent when they do. This means that, like the derivational opacity approach described above, underspecification theories of transparency are unable to account for skipping of harmonically paired vowels, as in Finnish loanwords.

For the medial [y:] in a word like [martty:ri] to be treated as transparent ([martty:ri-a], with a back suffix), it must be underlyingly unspecified along the front/back dimension. Otherwise, spreading backness from the initial [a] to the suffix vowel would run afoul of the inviolable line crossing prohibition. However, the distribution of [y] and [u] in Finnish — even in loanwords — is not predictable; the backness of a high rounded vowel is not determined by general (or even positional) well-formedness conditions, but is in fact abitrary. This means that the [y] in words like [martty:ri] must be underlyingly specified as front. There is a fundamental incompatibility of these two requirements, and it is impossible to account for the patterns of transparency in Finnish loanwords while maintaining an inviolable prohibition on crossed association lines.

The problem is not limited to Finnish; any language where otherwise contrastive vowels are treated as transparent is a problem for the underspecification theory of transparency. In Khalkha, for example, [i] is transparent to rounding harmony, but high vowels contrast for rounding. Kiparsky (1981) provides a successful and insightful analysis of cases like Hungarian, where contrast and participation are closely connected, but is ultimately unable to account for the broader typology of transparency and opacity.

6.2.3 Coarticulation as Participation

Benus (2005); Benus and Gafos (2007) preserve strict locality in an entirely different way — by considering transparent vowels to be participants, even on the surface. While transparent vowels do not undergo *categorical* phonological alternations, they do still undergo gradient phonetic coarticulation — Benus (2005); Benus and Gafos (2007) present experimental evidence that, in Hungarian, a transparent [i] that is flanked by back vowels is in fact less front than an [i] that is flanked by front vowels. Benus and Gafos argue that this gradient effect should be represented phonologically, and is sufficient to locally predict the backness vowel of a following suffix.

Their analysis is situated in Articulatory Phonology (Browman and Goldstein, 1986), and operates on gestures as phonological representations. In a word like [pɔpiːr], the input contains a back vowel gesture followed by a front vowel gesture. Gradient coarticulation and categorical harmony are driven by the same force: a requirement to minimize the distance between adjacent gestures. While [i:] is unable to fully comply by becoming back, the preceding back vowel exerts a coarticulatory influence, causing [i] to be realized as a slightly retracted front vowel. This is illustrated by the dashed line in (155).

Because [i:] is retracted, it can be considered to show some manifestation of that back gesture — backness is part of its representation, even though the vowel itself is still categorically front. In their view, this provides sufficient information to transmit backness to the suffix vowel — whether the suffix will be front or back is determined by whether or not [i:] manifests backness by being retracted. This is illustrated by the solid line in (155).

(155) pɔpi̠ːr-nɑk

The account is strictly local because the backness value of each vowel can be determined solely by considering information contained in the preceding vowel (the adjacent vowel gesture). Transparent vowels, while not categorical undergoers, are *transmitters* of backness, and therefore are considered to be participants.

Benus (2005) casts this analysis within an OT framework. Separate mechanisms are required to drive harmony from a categorically back vowel and transmission of backness from a retracted (but categorically front) vowel. The first, formalized in OT using the constraint AGREE(CL), is a requirement that adjacent segments minimize the distance between the loci of their articulatory gestures. This constraint is divided into both a root-specific, gradiently evaluated version (156a.i) and a suffix-specific, categorically-evaluated version (156a.ii). This constraint results in either harmony (in the case of paired vowels) or sub-phonemic coarticulation (in the case of non-participants [i] and [e]).

However, this constraint cannot be responsible for transparent harmony — despite the fact that [i] and [e] are less front than usual when they are preceded by a back vowel (gradiently satisfying AGREE(CL), it still involves less articulatory movement to select a *front* suffix in those contexts, but transparent harmony requires a back suffix. A second mechanism is used to drive harmony in these contexts; rather than operating on the distance between articulatory gestures, this constraint operates on *degree of retraction* — that is, the degree to which a front vowel (here [i] or [e]) is articulated with a less extreme gesture than it otherwise would be. If a front vowel is sufficiently retracted — that is, if gradient sub-phonemic coarticulation has caused it to be articulated with a sufficiently less extreme gesture — the constraint AGREE(R) requires that it be followed by a back vowel (156b).

(156) **Constraint definitions: Benus (2005)**

a. i. AGREE(CL)$_{\mathrm{St}}$: Consecutive stem vowels minimize their difference (distance) in terms of tongue body constriction location.

ii. Agree(CL)$_{\text{A-Suff}}$: Stem-final [+back] vowels minimize their articulatory distance with the suffix vowel(s) in terms of constriction location.

b. Agree(R)$_{\text{I-Suff}}$: Stem final [–back] vowels determine the quality of the suffix based on their retraction degree R.

Benus (2005) also distinguishes between articulatory (Ident(CL)) and perceptual faithfulness (Ident(front)). A segment is articulatorily unfaithful if its articulatory target in the output differs from its articulatory target in the input; this is gradiently evaluated, based on the distance between the two targets. A segment is perceptually unfaithful if its output realization will be perceived as belonging to a different category than its input specification; this is categorically evaluated.

The tableau in (157) demonstrates how these constraints are used to model transparent harmony in Hungarian. The input is a stem with an [a]–[i] vowel sequence, with an unspecified vowel in the dative suffix. The parameters for constriction location (CL), degree of retraction (R), and articulatory distance (AD) are schematized, with [a] given a CL of –2 and [i] given an input CL of +2. It is assumed that high-ranked initial syllable faithfulness accounts for directionality effects, and all candidates leave the parameters of [a] unchanged.

The medial [i] is preserved faithfully in (157ab) — both candidates incur the highest possible number of violations on Agree(CL)$_{\text{St}}$, but no violations on either faithfulness constraint.[2] In (157a), a back vowel is selected for the suffix — because the [i] in the stem is unretracted, this violates Agree(R)$_{\text{I-Suff}}$. In (157cd), the medial [i] has been backed categorically — this incurs a violation of the perceptual faithfulness constraint Ident(front) as well as two violations of the articulatory faith-

[2]Because Benus uses an unspecified suffix vowel, neither the candidate with a front suffix nor the candidate with a back suffix violate faithfulness. The same effect could be achieved with fully-specified suffix inputs so long as suffix-specific faithfulness constraints were sufficiently low-ranked.

fulness constraint Ident(CL), but it incurs the fewest violations of Agree(CL)$_{\text{St}}$. In (157d), a front suffix is selected; because the medial [i] is retracted to [ɯ], this incurs a violation of Agree(R)$_{\text{I-Stem}}$.

Finally, in (157ef), the medial [i] has been retracted, but to a lesser degree than in (157cd), resulting not in a back vowel [ɯ] but in a less extreme front vowel. This does not violate Ident(front), because it is still categorically front. Its violations of Agree(CL) are intermediate between those for the fully front vowel (157ab) and those for the fully back vowel (157cd); likewise for Ident(CL). In (157f), a front suffix has been selected — because the stem-final [i] is retracted (despite the fact that it is still categorically front), this incurs a violation of Agree(R)$_{\text{I-Suff}}$. The optimal candidate here is (157e) — it has orchestrated a compromise between Ident(front) and Agree(CL)$_{\text{St}}$ by retracting [i] enough to decrease the articulatory distance from [a], but not so much that it ceases to be categorically front. Because of this retraction, it must also satisfy Agree(R)$_{\text{I-Suff}}$ by selecting a back suffix.

(157) **Transparency in Hungarian**

(a i) v	Ident(fr)	Agr(R)$_{\text{I-Suff}}$	Agr(CL)$_{\text{St}}$	Ident(CL)
a. (a i) a		*!	***	
b. (a i) e			***!	
c. (a ɯ) a	*!		*	**
d. (a ɯ) e	*!	*	*	**
e. ☞ (a i̠) e			**	*
f. (a i̠) a		*!	**	*

Under this ranking, Benus (2005) is able to capture the behavior of transparent [i] and as well as the opacity of less perceptually stable front vowels. However, this ranking does not account for stem-internal harmony in native Hungarian words. This is because the constraint responsible for the fact that [i] does not *undergo* harmony is Ident(front) does not distinguish between [i] and vowels like [y] (which does undergo harmony). There are two possible ways around this issue; one is to either make the perceptual faithfulness constraint sensitive to the differences between [i] and other front vowels, or to substitute a markedness constraint like *[bk,rd,–lo] (which would be similarly sensitive to *category* membership, thus playing a similar role to the current IDENT(FRONT) in the tableau above). Another is to argue that this is in fact the correct interpretation of the facts in Hungarian, and that stem-internal harmony is no longer synchronically active, but rather an inherited fact about the lexicon — an argument of this type would be supported by the fact that recent loanwords aren't forced to harmonize (e.g. *parfüm* 'perfume').

Other problems with this approach, however, are not so easily resolved. In Chapter 5, I showed how an approach based on coarticulatory influence fails to explain the patterns seen in variable transparency and opacity in Finnish loanwords. With back–front nonce loanwords, subjects were more likely to choose transparency (a back suffix) when the initial vowel was low than when it was high — this is the precise opposite of what Benus (2005); Benus and Gafos (2007) predict. Furthermore, the theory advanced by Benus (2005); Benus and Gafos (2007) aims to account for the patterns of transparency and opacity in harmony proccess while maintaining strict locality — and, in doing so, providing a phonetically grounded account of the phenomenon. In Chapter 4, I showed how their approach does not in fact maintain phonetic grounding; the constraint AGREE(R) makes demands that are not justified phonetically. In attempting to preserve strict locality, it becomes necessary to depart from the very phonetic grounding motivating its preservation.

Related to this point about phonetic grounding, there is a phonological (and typological) argument to be made. The two harmony mechanisms, AGREE(CL) and AGREE(R), operate on clearly distinct parameters (articulatory distance and degree of retraction) and are functionally independent from on another. This gives rise to the possibility of a language Hungarian′, in which *only* AGREE(R) is active — that is, a language where a retracted [i] or [e] induces selection of a back suffix, but not

(a) a fully back vowel (as would be the case in a langauge with harmony) or (b) unretracted [i] or [e] (as would be the case in a language with vowel dissimilation).[3]

Benus (2005) avoids this problem by positing that suffixes in Hungarian are never specified for backness — this means that there is no faithfulness penalty for assimilating a suffix, and therefore nothing to prevent IDENT(CL)$_{\text{A-Suff}}$ from exerting its influence generally (wherever it does not conflict with AGREE(R)). However, in order to prevent Hungarian′ from ever being possible, it would be necessary to specify that suffixes could never be specified for backness, in any language. While underspecification may be a reasonable hypothesis in Hungarian (suffix values are always predictable), it is not tenable universally. Even languages which exhibit suffix alternations in harmony show contrasts in suffixes — for example, in exceptionally non-undergoing suffixes in Turkish (see 2.4.2 above). In addition, under Richness of the Base, even Hungarian would need to be equipped to handle the possibility of specified suffix vowels. Thus, a suffix-specific version of IDENT(front) has been included here, to reflect the possibility of languages with faithfulness to suffix feature values.

Hungarian′ is predicted under the ranking in (158). AGREE(CL)$_{\text{St}}$ dominates IDENT(CL), giving rise to coarticulatory retraction of a stem-final front vowel (additionally, IDENT(front)$_{\text{St}}$ dominates AGREE(CL)$_{\text{St}}$, ensuring that the result is a retracted front vowel rather than a categorically back vowel). AGREE(R)$_{\text{I-Suff}}$ dominates AGREE(CL)$_{\text{A-Suff}}$, meaning that following a retracted front vowel with a back suffix vowel is more important than the increase in articulatory distance introduced by such a configuration. However, IDENT(front)$_{\text{Suff}}$ dominates AGREE(CL)$_{\text{A-Suff}}$ — the mechanism involved in inducing suffix harmony from a back vowel is therefore unable to overcome perceptual faithfulness.

[3] In order for AGREE(R) to be relevant, AGREE(CL) must have an effect on a preceding vowel — however, because AGREE(CL) is divided into stem-specific and suffix-specific versions, it still remains possible to find a language where only AGREE(R) is able to induce suffix alternations of the kind seen in Hungarian.

(158) $\text{Agree(R)}_{\text{I-Suff}}$, $\text{Ident(front)}_{\text{St}}$ $\gg$ $\text{Agree(CL)}_{\text{St}}$ $\gg$ Ident(CL), $\text{Ident(front)}_{\text{Suff}}$ $\gg$ $\text{Agree(CL)}_{\text{A-Suff}}$

The tableaux in (159–161) demonstrate how this ranking produces Hungarian′. In (159), the input is a stem with a back–front vowel sequence and an underlyingly front suffix vowel. The set of possible candidates entertained is the same as in (157) above — front and back suffix forms with the medial vowel either surfacing faithfully (159ab), surfacing with coarticulatory retraction (159cd) or surfacing as fully back (159ef). Consider first the effects of Ident(front)$_{\text{St}}$, which prohibits categorical alternation of front vowels in stems. Candidates (159ef), which have categorically assimilated [i] to the preceding back vowel, incur fatal violations of this constraint, and will lose. Similarly, candidates (159ab), which have preserved the full frontness of [i], incur fatal violations of Agree(CL)$_{\text{St}}$, and will also lose.

The remaining candidates (159cd) have compromised between Agree(CL)$_{\text{St}}$ and Ident(front)$_{\text{St}}$, retracting the medial [i] but preserving its categorical frontness. The question now is the value of the suffix vowel; candidate (159c) surfaces with a faithful front vowel, satisfying the low-ranked Ident(front)$_{\text{Suff}}$ but incurring a fatal violation of the high-ranked Agree(R)$_{\text{I-Suff}}$. Its competitor, candidate (159d), has chosen a back suffix to satisfy Agree(R)$_{\text{I-Suff}}$, and is the winner. The transparent candidate here has won for the same reasons that it won in Benus's analysis shown in (157) above.

(159)

(a i) e	Agree (R)$_{\text{I-Suff}}$	Ident (front)$_{\text{St}}$	Agree (CL)$_{\text{St}}$	Ident (CL)	Ident (fr)$_{\text{Suff}}$	Agree (CL)$_{\text{A-Suff}}$
a. (a i) e			W2	L	L	
b. (a i) a	W1		W2	L	1	
c. (a i̠) e	W1		1	1	L	
d. ☞ (a i̠) a			1	1	1	
e. (a ɯ) e	W1	W1	L	W2	1	W2
f. (a ɯ) a		W1	L	W2	1	

However, the ranking for Hungarian′ produces unattested effects when other inputs are considered. In (160), the input consists of a stem with a back vowel and a suffix with a front vowel. Candidate (160a) is faithful to that front suffix, and incurs a maximal violation of AGREE(CL)$_{\text{A-Suff}}$; Candidate (160b) avoids that violation by selecting a front suffix, but loses because it violates the higher-ranked IDENT(front)$_{\text{Suff}}$. The winner here shows no harmony at all, despite a back stem vowel — underlyingly back stem vowels do not induce back suffix selection.

(160)

(a) e	AGREE (R)$_{\text{I-Suff}}$	IDENT (front)$_{\text{St}}$	AGREE (CL)$_{\text{St}}$	IDENT (CL)	IDENT (fr)$_{\text{Suff}}$	AGREE (CL)$_{\text{A-Suff}}$
a. ☞ (a) e						2
b. (a) a					W1	L

Furthermore, not all front vowels induce back suffix selection. In (161), the input has a front stem vowel and a front suffix vowel. Candidate (161a) surfaces faithfully, and does not violate any constraints; it harmonically bounds Candidate (161b), which surfaces with a back suffix, violating both AGREE(R)$_{\text{I-Suff}}$ and IDENT(front)$_{\text{Suff}}$.

(161)

(i) e	AGREE (R)$_{\text{I-Suff}}$	IDENT (front)$_{\text{St}}$	AGREE (CL)$_{\text{St}}$	IDENT (CL)	IDENT (fr)$_{\text{Suff}}$	AGREE (CL)$_{\text{A-Suff}}$
a. ☞ (i) e						
b. (i) a	W1				W1	

Thus, Hungarian′ is a language where an articulatorily retracted front vowel will induce back suffix selection, but front suffixes surface faithfully with both underlying back vowels and non-retracted front vowels. In a language like Hungarian′, then, harmony is *only* ever non-local. Furthermore, vowel-to-vowel coarticulation is not the only source of front-vowel retraction — [i] may be retracted due to coarticulatory influence from certain dorsal consonants, for example.

The reason a language like Hungarian′ is predicted to exist is that AGREE(CL) and AGREE(R) can't be put into a stringency relationship. They operate on distinct parameters — articulatory distance and degree of retraction — and therefore are predicted to behave independently. In the attested typology, there is an implicational relationship between local and non-local harmony — in languages where harmony can be non-local, *it can also be local.* Using functionally independent constraints to motivate local and non-local instances of harmony fails to capture this implicational relationship.

The competing triggers theory, on the other hand, does not separate the harmony mechanisms for local and non-local instances of harmony in this way. The pressure for locality is implemented as a diminution of the reward earned for harmony — meaning that local harmony will always earn more of a reward than non-local harmony, and therefore any language where non-local harmony is able to overcome faithfulness will necessarily be a language where local harmony is able to overcome faithfulness. A language like Hungarian′, therefore, is correctly predicted to be impossible.

6.2.4 The Case for Strict Locality

A common assumption in the literature on strict locality is that it provides a more phonetically grounded analysis of harmony (Archangeli and Pulleyblank, 1994; Pulleyblank, 2004; Gafos, 1999; Benus, 2005, and others). In Chapter 4, I provided evidence that this is not in fact the case. While strict locality is a necessary condition of phonetic grounding if articulatory ease is the sole motivating factor behind harmony, there are also *perceptual* advantages. Harmony increases the salience of a contrast by giving the listener multiple opportunities to recognize its cues, and this advantage is present both locally and non-locally. Additionally, because transparent harmony creates surface forms that are not strictly articulatorily local, attempting to maintain strictly local phonological interactions still involves a departure from phonetic grounding.

Another common argument in favor of strict locality is that it provides a simpler theory (Ni Chiosain and Padgett, 2001; Walker, 1998, and others). This sort of argument is virtually impossible to assess — comparing the parsimony of two theories requires that they make *identical predictions*. As we have seen above, approaches to transparency based on the assumption of strict locality make substantially different typological predictions and have different analytical capabilities than an approach based on explicitly non-local interactions. The simplicity argument is, at best, inapplicable.

Perhaps the most viable argument in favor of strict locality is the avoidance of seemingly contradictory precedence relationships. In a configuration like in (162), with association lines crossed, Feature A is associated with a segment that precedes Feature B *and* with a segment that follows it. Thus, Feature A both precedes and follows Feature B, making it impossible to determine a single linear order of A and B, and therefore making the structure impossible to realize.

(162) **Crossed Association Lines**

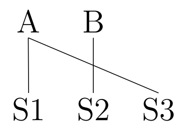

This argument is persuasive only if there is a direct, one-to-one mapping between autosegmental features and articulatory gestures. This is a central premise of approaches like Articulatory Phonology (Browman and Goldstein, 1986), but is far from being a universal and uncontroversial assumption. If autosegmental features are taken as phonological abstractions rather than as physical primitives, the contradictory precedence relationships involved are no longer a problem. The fact that A both precedes and follows B does not make the form unrealizable, because A and B are not themselves units which must be linearized — rather, it is the segments associated with those features which must (and can) be given a linear order.

Feature copying — rather than autosegmental associations — has been used to capture the effects of non-local harmony (Gafos, 1999; Hansson, 2001; Rose and Walker, 2004, and others). The problem of contradictory precedence relationships is avoided because there are, in fact, multiple (matching) features. One of the principles motivating the use of both autosegmental spreading (for local harmony) and feature copying (for non-local harmony) is the claim that there is a substantive difference between processes that can be local and processes that can be non-local.

Hansson (2001), for example, distinguishes consonant harmony processes, which are non-local and frequently similarity-sensitive, from vowel harmony processes. However, this distinction is not so tidy. Languages with transparent harmony (including detailed acoustic and articulatory evidence from Hungarian) show that vowel harmony can in fact be non-local; furthermore, while consonant harmony processes frequently operate at a distance, they are still subject to limitations and restrictions based on locality (see e.g. Odden 1994 for an overview). Additionally, languages with parasitic harmony show that consonant harmony is not alone in being similarity-sensitive (though see Section 6.3.2 for a discussion of how a theory based solely on similarity sensitivity fails to account for the attested patterns of transparency and opacity in vowel harmony).

Autosegmental relationships have been used to explain the positional or prosodic licensing of features (Goldsmith, 1989; Ringen, 1988, and many others) — a feature is licensed if one of the segments associated to is in a position that can host that feature. For example, Walker (2005) analyzes metaphony in various Romance languages (where high suffix vowels cause a stressed stem vowel to raise) as motivated by feature licensing. The feature [+HIGH] must be licensed by being associated with the prosodically prominent stressed syllable. However, Walker (2004) argues that feature copying relationships must also be sufficient for feature licensing, because metaphony may occur non-locally.

Another potential distinction between local and non-local harmony processes concerns multiple linking effects, such as Meussen's Rule, where a tone dissimilation process affects multiple segments associated with a single autosegment; a low tone causes a multiply-linked HH sequence to become LL (rather than LH, which would satisfy the requirements of dissimilation — see e.g. Myers 1997, and others). This kind of multiple-linking effect is frequently cited as evidence for autosegmental relationships — the fact that both high tones lower is only explained if they are the *same* high tone.

Archangeli and Pulleyblank (1994) cite the absence of multiple-linking effects in non-local assimilation as evidence in favor of strict locality — non-local assimilation processes, according to their argument, use feature copying rather than autosegmental spreading, and this is what explains the absence of those multiple linking effects. However, the absence of non-local multiple-linking effects can be equally well-explained by the combined rarity of both non-local harmony and multiple-linking effects. That is, we would need to find languages with (a) multiple-linking effects for adjacent assimilation, and (b) instances of non-adjacent assimilation.

To my knowledge, no such language happens to exist; none of the languages with multiple-linking effects in local harmony permit non-local harmony, and none of the langauges with non-local harmony show multiple linking effects. Because the prerequisites for the presence of non-local multiple-linking effects are not met, arguments based on its principled absence are not convincing.

In short, there is no substantive difference between the kinds of harmony processes that can be non-local and the kinds of processes that can't be — a feature copying theory that is sufficiently enriched to account for all the attested data on non-local harmony would become more or less indistinguishable from a theory based on autosegmental relationships. The difference between the two approaches is essentially notational rather than substantive, and the choice between them is primarily aesthetic. The proposal advanced in this dissertation could, with only slight modification, be adopted to rely on feature-matching relationships rather than autosegmental associations.

To summarize, strict locality doesn't provide a simpler or more phonetically grounded theory, and isn't representationally necessary. Furthermore, there is no fundamental difference between harmony that can be non-local and harmony that can't. There is, therefore, no *a priori* reason to prefer a strictly local account of harmony, and a theory that makes use of explictly non-local representation provides better empirical coverage.

6.3 Alternative Ways of Loosening Locality

6.3.1 Violable Locality

In OT-based frameworks, one way to relax the strictness of locality is to permit non-local representations, but introduce a markedness constraint penalizing non-local harmony. The precise definition of this constraint varies from proposal to proposal; I follow Walker (2010a) in adopting a generic version, Locality, as a stand-in (for specific formulations, see e.g. Cole and Kisseberth 1994; Kaun 1995; Rose and Walker 2004 and many others). Locality assigns violations to gapped configurations — that is, configurations where the segments associated with an autosegmental feature are non-contiguous. Here, Harmonize stands in for any harmony constraint that can apply non-locally.

A violable locality constraint can produce both transparent and opaque outcomes. If Harmonize dominates Locality, as in (163), satisfying the harmony imperative is more important and comes at the expense of violating the locality constraint — (163a), with transparent harmony, wins. On the other hand, if Locality dominates Harmonize, as in (164), the cost of non-locality is sufficiently great to block harmony — (164b), with opacity, is optimal instead.

(163)

analyyysi-ä	Harmonize	Loc	Ident
a. ☞ analyysi-a		1	1
b. analyysi-ä	W1	L	L

(164)

analyysi-ä	Loc	Harmonize	Ident
a. analyysi-a	W1	L	W1
b. ☞ analyysi-ä		1	

The problem with this approach is a too-many-solutions problem: there are other ways of satisfying the locality constraint besides blocking harmony. For example, if Harmonize and Locality dominate Max-V, as in (165), a candidate which has deleted intervening vowels will be the winner. The competition between harmony and locality has been rendered moot, because deletion *ensures* locality — (165c) satisfies both Locality and Harmonize, and emerges as the winner.

(165)

analyysi-ä	Loc	Harmonize	Max-V	Ident
a. analyysi-a	W1		L	1
b. analyysi-ä		W1	L	L
c. ☞ anals-a			2	1

This predicts a language where vowels delete whenever they are (a) ineligible for harmony, and (b) intervene between a trigger and a target[4] — no language does this.

The gradual derivations of Harmonic Serialism provide some relief, but do not entirely solve the problem. To begin with, harmony and deletion, being separate operations, cannot occur simultaneously; there is no candidate which has both deleted and spread. However, this does not preclude deletion as a solution to non-locality — with a ranking of Harmonize ≫ Locality, transparent harmony will be selected as optimal; subsequently, a ranking of Locality ≫ Max-V will force deletion to resolve the problem. However, if deletion must be gradual, as McCarthy (2008c) pro-

[4]Under some formulations of the harmony-driving constraint, this would result in ineligible vowels deleting whenever they co-occurred with a vowel bearing the spreading feature.

poses, then vowel deletion must pass through an intermediate step of vowel reduction. If locality is indeed governed by a single constraint, as shown above, this will provide no satisfaction — skipping a reduced vowel is no better than skipping a full vowel.

However, controlling locality via a single, omnibus constraint is not a viable approach. With only a single locality constraint, it is not possible to model quality-sensitive locality — so the asymmetries seen in Finnish loanwords are unexplained. For example, high vowels are more likely to be skipped than low vowels, meaning that there must be some rankings which produce transparent high vowels but opaque low vowels. This is impossible with a single locality constraint, as is shown in (166) — the ranking needed to produce transparent [y] and the ranking needed to produce opaque [ä] are contradictory, and cannot co-occur. No language is predicted to have both transparent and opaque segments.

(166)

winner ∼ loser	Harmonize	Loc	Ident
a. analyysi-a ∼ analyysi-ä	W	L	L
b. miljonääri-ä ∼ miljonääri-a	L	W	W

One solution to this problem, as proposed by Kaun (1995, and others) is to divide Locality into a family of quality-sensitive constraints; these can then be put into stringency relationships, or organized into a fixed ranking along a hierarchy (as proposed by Kaun 1995), to capture the cross-linguistic asymmetries in segments' conduciveness to transparency. Quality-sensitive locality constraints are able to capture languages with both transparent and opaque segments — in Finnish, for example, the ranking needed to get the transparency of [y] (Harmonize ≫ *Skip(high) is compatible with the ranking needed to get the opacity of [ä] (*Skip(low) ≫ Harmonize). This is shown in (167).

(167)

winner ∼ loser	*Skip(lo)	Harmonize	*Skip(hi)	Id
a. analyysi-a ∼ analyysi-ä		W	L	L
b. miljonääri-ä ∼ miljonääri-a	W	L		W

However, quality-sensitive locality constraints mean that Harmonic Serialism does not provide relief from the too-many-solutions problem — while vowel deletion can be rendered sufficiently gradual, by requiring the intermediate step of reduction, the same cannot necessarily be said about changes in vowel quality. A family of violable locality constraints predicts that, once transparent harmony has occurred, a skipped segment can then change along some other feature dimension to shift its locality-related violations to a lower-ranked constraint. This is shown in the derivation in (168).

At the first step of the derivation, the faithful candidate (a) incurs a violation of HARMONIZE, as does the candidate that has gratuitously raised the medial vowel (c). The winning candidate (b) harmonizes despite a violation of *SKIP(LOW) (and an additional violation of *SKIP(HIGH) for skipping the [i] as well). At the second step in the derivation, the faithful candidate (a) maintains its violation of *SKIP(LOW); however, the winning candidate (b) raises the medial vowel — it no longer incurs a violation of *SKIP(LOW), but rather an additional violation of *SKIP(HIGH), at the cost of violating the low-ranked IDENT(LOW).

(168) **Step 1**

	miljonääri-ä	HARM	*SKIP(LO)	*SKIP(HI)	ID(FRONT)	ID(LO)
a.	miljonääri-ä	W1	L	L	L	
b. ☞	miljonääri-a		1	1	1	
c.	miljonyyri-ä	W1	L	L	L	W1

Step 2

	miljonääri-a	HARM	*SKIP(LO)	*SKIP(HI)	ID(FRONT)	ID(LO)
a.	miljonääri-a		W1	L1		
b. ☞	miljonyyri-a			2		W1

The prediction here is a language where [a] raises to [y] just in case it intervenes between a harmony trigger and target (the [a] in the above example would surface faithfully in, for example, unsuffixed forms or with suffixes which were ineligible to undergo harmony). Violable locality makes these kinds of aberrant typological predictions because there are ways of satisfying a locality constraint besides failing to harmonize. The approach presented in this dissertation avoids these too-many-solutions problems because locality is not enforced by an independent constraint, but rather a scaling factor on the constraints motivating harmony.

The reason the competing triggers approach avoids transparent vowels changing along other dimensions is that the harmony constraint is positive — difference between a preferred trigger and a dispreferred trigger comes into play *only if that vowel triggers harmony*. For example, consider the Finnish data from Chapter 5 — if transparent harmony is chosen with an input stem like [tafsä], could a subsequent step raise [ä] to make it more skippable, as predicted above with violable locality?

In the tableau in (169), we see that the answer is no. The potential of [ä] to block harmony came from its power to trigger harmony itself, not from an a priori penalty for skipping a low vowel. Raising it to [y], as in candidate (169b), provides no benefit and gratuitously violates faithfulness; the faithful (169a) wins, and the derivation converges.

(169) **Step 2**: *Convergence*

tafsä-ssa	SPREAD(B) 3.297	SPREAD(F) 1.924	IDENT(LOW) 1	$\mathcal{H}$
a. ☞ tafsä-ssa	+0.375 (+1 * 1.5 * .25)			1.236
b. tafsy-ssa	+0.375 (+1 * 1.5 * .25)		-1	0.236

Note that had the spreading constraint been negatively defined, assigning more penalty for failing to spread from a low vowel than for failing to spread from a high vowel, a change in quality would be motivated — the violation for failing to spread would persist, and given that spreading from that segment has lost, changing to a poorer trigger would provide some relief from the constraint motivating spreading.

As discussed in 3.5 above, the scaling factor for distance would, at first, appear to result in a similar too-many-solutions problem, this time not avoided by the positive formulation of the constraint. Consider again the Finnish data from Chapter 5; if harmony across a long vowel is chosen for a word like [paftyy], will shortening the skipped vowel mean that the instance of harmony is subject to the less severe scaling factor for distance?

The answer, again, is no — recall that the scaling factors are a property of *operations* rather than a property of representations. The scaling for non-local harmony applies for determining an instance of assimilation; at the step where the harmony operation occurred, a long vowel intervened between target and trigger. Because there is no longer an active trigger once the operation has been completed, the scaling factors are no longer subject to adjustment.

What this means is that, if transparent harmony across a long vowel is chosen at the first step of a derivation, there will be no advantage to shortening the vowel subsequently. This is shown in the tableau in (170) — because the initial [a] is no longer a *trigger*, the scaling factors are no longer subject to assessment. The candidates which preserve the association link in the input will preserve only the static information about how its reward should be scaled. The rewards on SPREAD(B) are the same for both the faithful candidate (170b) and the candidate with a shortened medial vowel (170a), because they are derived from the same input; shortening that vowel gratuitiously violates faithfulness, and the faithful candidate emerges as the winner (and the derivation converges).

(170) **Step 2**: *Convergence*

paftyy-sta	SPREAD(B) 3.297	SPREAD(F) 1.924	ID(LONG) 1	$\mathcal{H}$
a. pafty-sta	+0.233 (+1 * 1.5 * .155)		-1	-0.232
b. ☞ paftyy-sta	+0.233 (+1 * 1.5 * .155)			0.767

To summarize, the competing triggers account proposed here avoids the too-many-solutions problems suffered by violable locality constraints because (a) the harmony-driving constraints are positive, assigning rewards rather than violations, and (b) the scaling factor for locality is assessed over operations rather than representations, and is therefore not subject to subsequent adjustment.

6.3.2 Agreement By Correspondence

The theory of Agreement By Correspondence (ABC), designed to handle non-locality in consonant harmony, capitalizes on the claim that only *irrelevant* segments are skipped. Harmony is achieved via surface correspondence among segments; segments which share a particular feature are compelled to stand in a correspondence relationship (by CORR-SS constraints), and segments which stand in correspondence are compelled to agree with respect to the harmonizing feature (by IDENT-SS constraints).

Hansson (2001); Rose and Walker (2004) show how this provides a satisfactory account of consonant harmony processes — in particular, the basis on which segments are compelled to stand in correspondence captures the generalization that consonant harmony tends to be highly sensitive to similarity. The account can be extended to analyze non-locality in vowel harmony, but only insofar as vowel harmony is also similarity-sensitive.

For example, in Yoruba, the segments which participate in harmony can be classified along a single feature dimension that excludes the segments which do not participate — namely, mid vowels undergo (right-to-left) RTR harmony, and other vowels (in particular, high vowels) do not. To capture this generalization, CORR-OE requires mid vowels to stand in correspondence with each other, and IDENT-SS(RTR) requires (chain-adjacent) corresponding segments to agree with respect to RTR.

The ranking in (171) produces transparency of high vowels, as in the Ifẹ dialect of Yoruba. In candidates (171a–d), all vowels stand in correspondence. The faithful candidate (171a) incurs a violation of IDENT-SS(RTR), because the medial [u] and the final [ɔ] disagree with respect to RTR. Candidate (171b) assimilates the initial [e] to [ɛ], and incurs two violations of IDENT-SS(RTR) (one for the disagreement between [ɛ] and [u], and one for the disagreement between [u] and [ɔ]). Candidates (171cd) both change the medial [u] to [ʊ], incurring violations of the feature co-occurrence constraint *HI/RTR.

In candidate (171f), the final [ɔ] is left out of correspondence; this means that no violations of IDENT-SS(RTR) are incurred (cf. 171a), but CORR-OE is violated — [e] and [ɔ] are both mid vowels, and are not in correspondence with each other. The winner is candidate (171e), which has left the medial [u] out of correspondence (because [u] is not mid, this does not violate CORR-OE), and changed the initial [e] to [ɛ] to agree with [ɔ] (avoiding a violation of IDENT-SS(RTR)).

(171)

elubɔ	CORR-OE	*HI/RTR	IDENT-SS(RTR)	IDENT-IO(RTR)
a. $\text{e}_i\text{lu}_i\text{bɔ}_i$			W1	L
b. $\text{ɛ}_i\text{lu}_i\text{bɔ}_i$			W2	1
c. $\text{e}_i\text{lʊ}_i\text{bɔ}_i$		W1	W1	1
d. $\text{ɛ}_i\text{lʊ}_i\text{bɔ}_i$		W1		W2
e. ☞ $\text{ɛ}_i\text{lu}_j\text{bɔ}_i$				1
f. $\text{e}_i\text{lu}_i\text{bɔ}_j$	W1			L

In (171), the [u] is skippable because it is irrelevant for assessing the constraint compelling correspondence. Because of this, it was originally thought that blocking was impossible under ABC, however Hansson (2007) notes a circumstance under which opacity is predicted — when a segment meets the criteria to be compelled to correspond, but is unable to participate. This is demonstrated (for the Ọyọ dialect of Yoruba) in (172), where the constraint CORR-VV is used instead of CORR-OE; CORR-VV compells all vowels to stand in correspondence with each other.

The same range of candidates are considered in (172) as were considered in (171). This time, however, candidate (172e) does not escape CORR-VV, and is compelled into correspondence. Candidates (172c–d) are once again ruled out by the high-ranking feature co-occurrence constraint. The choice is between the faithful candidate (172a) and candidate (172b), which has changed [e] to [ɛ] to agree with [ɔ]. Candidate (172b) incurs two violations of IDENT-SS(RTR) — one for the disagreement between [ɛ] and [u], and one for the disagreement between [u] and [ɔ]. Candidate (172a) only incurs one violation (for disagreement between [u] and [ɔ]), and is the winner.

(172)

elubɔ	CORR-VV	*HI/RTR	IDENT-SS(RTR)	IDENT-IO(RTR)
a. ☞ $\text{e}_i\text{lu}_i\text{bɔ}_i$			1	
b. $\text{ɛ}_i\text{lu}_i\text{bɔ}_i$			W2	W1
c. $\text{e}_i\text{lʊ}_i\text{bɔ}_i$		W1	1	W1
d. $\text{ɛ}_i\text{lʊ}_i\text{bɔ}_i$		W1	L	W2
e. $\text{ɛ}_i\text{lu}_j\text{bɔ}_i$	W2		L	W1
f. $\text{e}_i\text{lu}_i\text{bɔ}_j$	W2		L	

The tableaux in (171) and (172) show that ABC can produce transparency when the participating segments share some feature to the exclusion of the transparent segment, and can produce opacity when the opaque segment is sufficiently similar to the participating segments to be compelled into correspondence, but unable to undergo harmony. Furthermore, Walker (2009) demonstrates that ABC can produce a language with both transparency and opacity — when the opaque segment is similar to participating segments, but the transparent segment is not.

However, ABC enounters problems when attempting to account for the full range of typological patterns. First, even for consonant harmony processes, different *degrees* of non-locality must be taken into consideration (see e.g. Odden 1994 for examples) — but ABC treats transparent segments as though they don't exist, and so provides no mechanism for determining distance. This requires including some sort of locality constraint, such as Proximity (Rose and Walker, 2004), reintroducing the problems with violable locality seen above.

The other problem is that, while vowel harmony is sometimes similarity-sensitive (as in parasitic harmony), it is not *always* similarity sensitive. ABC works well in accounting for languages where the set of participants can be characterized by a feature or features that specifically excludes the set of non-participants. However, this is not always the case.

Consider, for example, the set of participating (and non-participating) segments in Finnish — in the native phonology, participants are [y,ø,æ,u,o,a] and non-participants are [i,e]. In order for [i,e] to be transparent, there must be a feature or features that the participants share, but [i,e] do not. There is no such feature; there is nothing that can be used to identify *all and only* participating vowels. Appealing to multiple features can characterize all and only *non*–participants, but what is required is a criterion for inclusion, not exclusion.

To get around this problem, Rhodes (2010) proposes an additional feature — the property of being contrastive in the inventory. Setting aside the fact that this approach requires a conception of the inventory as a linguistic object that can be referenced by constraints, it still is not true that this additional feature can explain the full range of possible transparent segments. In Finnish, for example, front vowels [y,ø,æ] can be optionally transparent in disharmonic loanwords; furthermore, they are not equally likely to be treated as transparent. To get the asymmetries in transparency, there must be possible rankings where [y] is transparent (along with [i,e]), but [ø] and [æ] are not — even with the help of a new feature like ±contrastive, there is no way to characterize the set of participants.

It should be noted that this is not a problem with segmental correspondence itself; the difference between autosegmental associations and segmental correspondence relationships is primarily one of notation, not substance. Rather this problem arises from the mechanisms for inducing correspondence, and the mechanism for enforcing agreement among corresponding segments. In other words, the problem arises as a result of the particular constraint set used by ABC to govern locality in harmony.

The competing triggers theory does not suffer from this problem because there is no need to cohesively unite the set of participating segments. It is still possible to capture instances of similarity sensitivity — a constraint like Kaun (1995)'s Uniformity would be compatible with the proposal presented in this dissertation, as would a scaling factor based on identity of trigger and target. In short, an approach like ABC is not required to achieve some degree of similarity sensitivity, and the competing triggers approach is also able to provide an analysis of transparency in cases where non-participation is not based on similarity.

6.4 Conclusion

In this chapter, I have compared the competing triggers approach proposed in this dissertation with a number of previous approaches to the problem of transparency and opacity in vowel harmony.

I showed that analyses based in strict locality ultimately require abandoning the phonetic grounding that preserving strict locality aims to maintain. Furthermore, they fail to account for the full range of data — derivational opacity fails to account for optional transparency of paired vowels in Finnish loanwords; coarticulation-as-participation predicts effects in the wrong direction for initial-vowel effects in the Finnish experiment from Chapter 5; underspecification fails to make any predictions about cross-linguistic asymmetries in segments' amenability to transparency. Additionally, Benus (2005)'s analysis incorrectly predicts a possible language where harmony is *only* ever non-local.

Extant approaches which depart from strict locality also encounter difficulties. Using a violable constraint to exert a preference for locality is plagued by a too-many-solutions problem — segments could delete (or change along some other feature dimension) to satisfy both the harmony imperative and the locality constraint. Agreement by Correspondence, on the other hand, while successful at modeling consonant harmony systems, fails to account for vowel harmony systems in which assimilation is not parasitic.

I have shown how the competing triggers analysis avoids each of these problems, providing better empirical coverage than previous accounts of transparency and opacity.

CHAPTER 7
CONCLUSION

In this dissertation, I have argued for a theory of transparency and opacity in vowel harmony in which the choice between the two options represents a competition between potential harmony triggers — the dominant but non-local trigger on the one hand and the non-dominant but local trigger on the other. Quality-based differences in trigger strength can influence this choice — segments are more likely to be transparent when they are poor triggers, and more likely to be opaque when they are strong triggers.

The proposal advanced here makes use of explicitly non-local representations, arguing that transparency does in fact involve phonological interaction between segments which are not adjacent to one another. The question of strict locality, and whether such representations should be permitted in the first place, is given careful examination. I address both the functional and formal arguments favoring maintaining strict locality, and I conclude that an absolute inviolable locality restriction is neither necessary nor tenable. I present experimental evidence showing that harmony among non-local segments confers the same perceptual benefit as harmony among local segments — transparent harmony, rather than being in some way exceptional or unnatural, is a natural phenomenon as part of a perceptually grounded theory of harmony. Additionally, I show that a theory which permits non-local representations provides better empirical coverage than a theory which does not.

Including trigger strength as an influence on transparency and opacity is supported by both experimental and typological data. I show that the same properties which influence asymmetries in trigger preference cross-linguistically are those which separate opaque non-participants from transparent non-participants in languages where transparency and opacity coexist. In particular, segments whose cues with respect to the harmonizing feature are less salient are both preferred as triggers independently and are more likely to be treated as opaque in the case of their non-participation. This is a prediction which follows from the broader claim, supported with experimental evidence, that harmony is perceptually motivated. Furthermore, the competing triggers theory predicts that the properties of both potential triggers should influence the outcome in this way — a prediction which is supported by results experimental evidence from Finnish disharmonic loanwords.

Finally, I have argued that factors which influence trigger strength should be formally represented as scaling factors on a positively-formulated harmony constraint, an analysis situated in the framework of Serial Harmonic Grammar. I argue that the combination of a positive constraint and serial evaluation provides a better overall account of harmony than previous alternatives, and that cumulative constraint interaction enables scalar constraints to account for the kind of hierarchies and asymmetries seen in preferential trigger effects. Furthermore, I show that defining scaling factors operationally rather than representationally ensures that preferences for locality and strong triggers are limited in their power — they are able to prevent harmony, but not induce other changes.

In conclusion, the proposal advanced in this dissertation provides an account of transparency and opacity in harmony which is justified both in its formal predictions and in its functional underpinnings, and supported with both typological and experimental evidence.

BIBLIOGRAPHY

Akinlabi, Akinbiyi. 2009. Neutral vowels in Lokaa harmony. *The Canadian Journal of Linguistics* 197–228.

Alderete, John. 1997. Dissimilation as local conjunction. In *Proceedings of the Northeast Linguistics Society*, ed. Kiyomi Kusumoto, volume 27, 17–32.

Archangeli, Diana. 1984. Underspecification in Yawelmani phonology and morphology. Doctoral Dissertation, MIT.

Archangeli, Diana, and Douglas Pulleyblank. 1994. *Grounded phonology*. MIT Press.

Archangeli, Diana, and Douglas Pulleyblank. 2002. Kinande vowel harmony: Domains, grounded conditions, and one-sided alignment. *Phonology* 19:139–188. University of Arizona and University of British Columbia.

Baayen, R. H. 2010. *languageR: Data sets and functions with "Analyzing Linguistic Data: A practical introduction to statistics"*. URL `http://CRAN.R-project.org/package=languageR`, R package version 1.0.

Bach, E. 1968. Two proposals concerning the simplicity metric in phonology. *Glossa* 2:128–149.

Bakovic, Eric. 2000. Harmony, dominance, and control. Doctoral Dissertation, University of California, Santa Cruz. ROA-360.

Bakovic, Eric, and Colin Wilson. 2000. Transparency, strict locality, and targeted constraints. In *WCCFL 19 Proceedings*, ed. R. Billerey and B. Lillehaugen.

Bates, Douglas, and Martin Maechler. 2010. *lme4: Linear mixed-effects models using s4 classes*. URL `http://CRAN.R-project.org/package=lme4`, R package version 0.999375-33.

Beckman, Jill. 1998. Positional faithfulness. Doctoral Dissertation, University of Massachusetts, Amherst.

Benus, Stefan. 2005. Dynamics and transparency in vowel harmony. Doctoral Dissertation, New York University.

Benus, Stefan, and Adamantios I. Gafos. 2007. Articulatory characteristics of Hungarian transparent vowels. *Journal of Phonetics* 35.

Bloomfield, Leonard. 1962. *The Menomini language*. Yale University Press.

Boersma, Paul. 1998a. Functional phonology: Formalizing the interactions beteween articulatory and perceptual drives. Doctoral Dissertation, University of Amsterdam.

Boersma, Paul. 1998b. Spreading in functional phonology. Ms, University of Amsterdam.

Boersma, Paul, and Bruce Hayes. 2001. Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32:45–86.

Boersma, Paul, and Joe Pater. 2008. Convergence properties of a gradual learning algorithm for Harmonic Grammar. ROA 970.

Boersma, Paul, and David Weenink. 2008. Praat: doing phonetics by computer (Version 5.0.17) [Computer program]. Retrieved April 1, 2008 from http://www.praat.org/. Developed at the Institute of Phonetic Sciences, University of Amsterdam.

Boyce, S. E. 1990. Coarticulatory organization for lip rounding in Turkish and English. *Journal of the Acoustical Society of America* 88:2584–2595.

Browman, C. P., and L. Goldstein. 1986. Towards and articulatory phonology. *Phonology Yearbook* 3:219–252.

Cammenga, Jillert. 1994. Kuria phonology and morphology. Doctoral Dissertation, Vrije Universiteit.

Campbell, Lyle. 1980. The psychological and sociological reality of Finnish vowel harmony. In *Issues in vowel harmony*, ed. Robert M. Vago, 245–269. John Benjamins.

Casali, Roderic. 2008. ATR harmony in African languages. *Language and Linguistic Compass* 2:496–549.

Clements, G. N. 1976. Neutral vowels in Hungarian vowel harmony: an autosegmental interpretation. In *NELS 7*, 49–64.

Clements, G. N., and Engin Sezer. 1982. Vowel and consonant disharmony in turkish. In *The structure of phonological representations*, ed. Harry van der Hulst and Norval Smith. Dordrecht: Foris Publications.

Clements, George N. 1977. The autosegmental treatment of vowel harmony. In *Phonologica 1976*, ed. W. Dressler and O. Pfeiffer.

Clements, George N. 1990. The hierarchical representation of vowel height. Ms, Cornell University.

Clements, George N. 1991. Vowel height assimilation in Bantu languages. *Working papers of the Cornell Phonetics Laboratory* 5:37–76.

Coetzee, Andries. 2009. An integrated grammatical/non-grammatical model of phonological variation. In *Current issues in linguistic interfaces*, ed. Young-Se Kang, Yong-Yurl Yoon, Hyunkyung Yoo, Sze-Wing Tang, Yong-Soon Kang, Youngjun Jang, Chul Kim, Kyoung-Ae Kim, and Hye-Kyung Kang, volume 2, 267–294. Hankookmunhwasa.

Coetzee, Andries, and Joe Pater. 2008. The place of variation in phonological theory. ROA 946.

Cole, Jennifer. 1987. Planar phonology and morphology. Doctoral Dissertation, MIT.

Cole, Jennifer, and Charles W. Kisseberth. 1994. An optimal domains theory of harmony. Urbana: University of Illinois. ROA-22.

Cole, Jennifer, and Charles W. Kisseberth. 1995a. Nasal harmony in Optimal Domains Theory. Ms., University of Illinois. ROA-49.

Cole, Jennifer, and Charles W. Kisseberth. 1995b. Restricting multi-level constraint evaluation: Opaque rule interaction in Yawelmani vowel harmony. Ms, University of Illinois, Urbana.

Cole, Jennifer, and L. Trigo. 1987. On the representation of neutral segments in harmony systems. Ms., MIT.

Cole, Jennifer, and L. Trigo. 1988. Parasitic harmony. In *Features, segmental structure, and harmony processes ii*, ed. Harry van der Hulst and Norval Smith, 19–38. Foris.

Cook, Eung-Do. 1976. Flattening and rounding in Chilcotin velars. In *Victoria Conference on Northwestern Languages*, ed. Barbara S. Efrat, 15–32.

Donnelly, Simon. 2000. Innovation in Tekela Nguni: Superclose and tense/lax vowel harmonies in Phuthi. Paper presented at LSSA 2000, Cape Town.

Eisner, Jason. 1999. Doing OT in a straightjacket. Colloquium, University of California, Los Angeles.

Elfner, Emily. 2009. Syllabification and stress-epenthesis interactions in Harmonic Serialism. ROA 1047.

Finley, Sara. 2008. Formal and cognitive restrictions on vowel harmony. Doctoral Dissertation, Johns Hopkins University.

Finley, Sara. 2009. Locality restrictions on exceptions to vowel harmony. *University of Rochester Working Papers in the Language Sciences* 5:17–48.

Flemming, Edward. 2001. Scalar and categorical phenomena in a unified model of phonetics and phonology. *Phonology* 18.

Gafos, Adamantios I. 1998. Eliminating long-distance consonantal spreading. *Natural Language and Linguistc Theory* 223–278.

Gafos, Adamantios I. 1999. The articulatory basis of locality in phonology. Doctoral Dissertation, Johns Hopkins University.

Gallagher, Gillian. 2010. The perceptual basis of long-distance laryngeal restrictions. Doctoral Dissertation, MIT.

Gick, Bryan, Douglas Pulleyblank, Fiona Campbell, and Ngessimo Mutaka. 2006. Low vowels and transparency in Kinande vowel harmony. *Phonology* 23:1–20.

Goldsmith, J. 1985. Vowel harmony in Khalkha Mongolian, Yaka, Finnish, and Hungarian. *Phonology Yearbook* 2.

Goldsmith, John. 1989. Autosegmental licensing, inalterability, and harmonic rule application. In *Papers from the 25th Meeting of the Chicago Linguistics Soceity*, 145–156.

Goldwater, Sharon, and Mark Johnson. 2003. Learning OT constraint rankings using a maximum entropy model. In *Proceedings of the Workshop on Variation within Optimality Theory*, ed. Jennifer Spenader, Anders Eriksson, and Östen Dahl, 111–120. Stockholm University.

Gordon, Matthew K. 1999. The "neutral" vowels of Finnish: How neutral are they? *Linguistica Uralica* 35:17–21.

Gouskova, Maria. 2003. Deriving economy: Syncope in optimality theory. Doctoral Dissertation, University of Massachusetts, Amherst.

Haiman, John. 1972. Phonological targets and unmarked structures. *Language* 48:365–377.

Hall, Beatrice L., and R.M.R. Hall. 1980. Nez Perce vowel harmony: An Africanist explanation and some theoretical questions. In *Issues in vowel harmony*, ed. Robert M. Vago. Benjamins.

Hansson, Gunnar. 2001. Theoretical and typological issues in consonant harmony. Doctoral Dissertation, University of California, Berkeley.

Hansson, Gunnar. 2007. Blocking effects in Agreement by Correspondence. *Linguistic Inquiry* 38:395–409.

Hare, Mary. 1990. The role of similarity in Hungarian vowel harmony: a connectionist account. *Connection Science* .

Hayes, Bruce, and Zsuzsa Cziráky Londe. 2006. Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology* 23:59–104.

Hayes, Bruce, Kie Zuraw, Péter Siptár, and Zsuzsa Cziráky Londe. 2009. Natural and unnatural constraints in Hungarian. *Language* 85:822–863.

Hetzron, Robert, and Habte-Miriam Marcos. 1966. Des traits pertinents superposés en ennemor. *Journal of Ethiopian Studies* 4:17–30.

Hyman, Larry. 1989. Advanced tongue root in Kinande. Ms, Unviersity of California, Berkeley.

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

Jäger, Gerhard, and Anette Rosenbach. 2006. Maximum entropy models and stochastic Optimality Theory. *Linguistics* 44:937–971.

Jensen, J. T. 1974. A constraint on variables in phonology. *Language* 50:675–686.

Jesney, Karen. 2009. Positional faithfulness, non-locality, and the Harmonic Serialism solution. In *Proceedings of the 39th Meeting of the North East Linguistics Society (NELS 39)*, ed. Suzi Lima, Kevin Mullin, and Brian Smith.

Jesney, Karen. to appear. Licensing in multiple contexts: an argument for Harmonic Grammar. In *Proceedings of the 45th Meeting of the Chicago Linguistics Soceity*.

Jurgec, Peter. 2011. Feature spreading 2.0: A unified theory of assimilation. Doctoral Dissertation, University of Tromsø.

Kaun, Abigail. 1995. The typology of rounding harmony: An Optimality Theoretic approach. Doctoral Dissertation, University of California, Los Angeles.

Kaye, Jonathan. 1982. Harmony processes in Vata. In *Projet sur les langues Kru*, ed. Jonathan Kaye, Hilda Koopman, and Dominique Sportiche, 60–151. Montreal, Canada: University of Quebec at Montreal.

Kiparsky, Paul. 1973. Phonological representations. In *Three dimensions in linguistic theory*, ed. O. Fujimura, 2–136. TEC.

Kiparsky, Paul. 1981. Vowel harmony. MIT.

Kiparsky, Paul. 1985. Some consequences of Lexical Phonology. *Phonology Yearbook* 2:83–138.

Kiparsky, Paul, and Karl Pajusalu. 2003. Towards a typology of disharmony. *The Linguistic Review* 20:217–241.

Kirchner, Robert. 1993. Turkish vowel harmony and disharmony: An Optimality Theoretic account. Ms., UCLA. ROA-4.

Kontra, Miklós, and Catherine Ringen. 1986. Hungarian vowel harmony: The evidence from loanwords. *Ural-Altaische Jahrbucher* 58:1–14.

Kornai, András. 1991. Hungarian vowel harmony. In *Approaches to Hungarian III*, ed. István Kenesei, 183–240.

Krämer, Martin. 2003. *Vowel harmony and correspondence theory*. Mouton de Gruyter.

Krauss, Michael E. 1975. Chilcotin phonology: a descriptive and historical report, with recommendations for a Chilcotin orthography. Alaska Native Language Center.

Krueger, John. 1962. *Yakut manual*. Indiana University Press.

Kullback, S, and R.A. Leibler. 1951. On information and sufficiency. *Annals of Mathematical Statistics* 22:79–86.

Lewis, Geoffrey L. 1967. *Turkish grammar*. Oxford University Press.

Li, Bing. 1996. *Tungusic vowel harmony*. Holland Institute of Generative Linguistics.

Lindau-Webb, Mona. 1987. Acoustic correlates of TENSE-LAX and ATR vowels. *Journal of the Acoustic Society of America* 82:S116–S116.

Linker, Wendy. 1982. Articulatory and acoustic correlates of labial activity in vowels: A cross-linguistic study. Doctoral Dissertation, UCLA.

Lombardi, Linda. 1999. Positional faithfulness and voicing assimilation in Optimality Theory. *Natural Language and Linguistic Theory* 17:267–302.

Lombardi, Linda, ed. 2001. *Segmental Phonology in Optimality Theory. Constraints and Representations*. Cambridge University Press.

Lynch, John. 2003. Low vowel dissimilation in Vanuatu languages. *Oceanic Linguistics* 42:359–406.

McCarthy, John. 2000. Harmonic Serialism and parallelism. In *Proceedings of NELS 30*, ed. Masako Hirotani, 501–24. Amherst, MA: GLSA Publications.

McCarthy, John. 2003a. Comparative markedness. *Theoretical Linguistics* 29:1–51.

McCarthy, John. 2003b. Sympathy, cumulativity, and the Duke-of-York gambit. In *The syllable in Optimality Theory*, 23–76. Cambridge University Press.

McCarthy, John. 2004. Headed spans and autosegmental spreading. ROA-695.

McCarthy, John. 2007a. *Hidden generalizations: Phonological opacity in optimality theory*. London: Equinox Publishing.

McCarthy, John. 2007b. Restraint of analysis. In *Freedom of analysis*, ed. Sylvia Blaho, Patrick Bye, and Martin Krämer, 203–31. Berlin and New York: Mouton de Gruyter.

McCarthy, John. 2008a. The gradual path to cluster simplification. *Phonology* 25:271–319.

McCarthy, John. 2008b. The serial interaction of stress and syncope. *Natural Language and Linguistc Theory* 26:499–546.

McCarthy, John. 2008c. Son of span: Harmony in Harmonic Serialism. University of Massachusetts, Amherst.

McCarthy, John. 2010. Studying Gen. *Journal of the Phonetic Society of Japan* .

McCarthy, John. to appear. Autosegmental spreading in Optimality Theory. In *Tones and features*, ed. John Goldsmith, Elizabeth Hume, and Leo Wetzels. Mouton de Gruyter.

Milligan, Marianne. 2000. A new look at Menominee vowel harmony. In *Papers of the 31st Algonguian Conference*, 237–254.

Miner, Kenneth L. 1975. *Omaeqnomenew-kiketwanan: a basic English-Menominee and Menominee-English word list.*. Wisconsin Department of Public Instruction.

Mullin, Kevin. 2010. Strength in harmony systems: Trigger and directional asymmetries. Ms., University of Massachusetts, Amherst.

Mutaka, Ngessimo. 1995. Vowel harmony in Kinande. *Journal of West African Languages* 25:41–55.

Myers, Scott. 1997. OCP effects in optimality theory. *Natural Language & Linguistic Theory* 15:847–92.

Nevins, Andrew. 2004. Conditions on (dis)harmony. Doctoral Dissertation, Massachusetts Institute of Technology.

Ni Chiosain, Maire, and Jaye Padgett. 2001. Markedness, segment realization, and locality in spreading. In *Segmental phonology in optimality theory: Constraints and representations*, ed. Linda Lombardi. New York: Cambridge University Press.

Noske, Manuela. 1996. [ATR] harmony in Turkana. *Studies in African Linguistics* 25:61–99.

Odden, David. 1991. Vowel geometry. *Phonology* 8:261–289.

Odden, David. 1994. Adjacency parameters in phonology. *Language* 70:289–330.

Onn, Farid. 1980. Aspects of Malay phonology and morphology. Doctoral Dissertation, Universiti Kebangsaan Malaysia.

Orie, Ọlanike Ọla. 2001. An alignment-based account of vowel harmony in Ifẹ yoruba. *Journal of African Languages and Linguistics* 22:117–143.

Orie, Ọlanike Ọla. 2003. Two harmony theories and high vowel variability in two Benue-Congo languages. *The Linguistic Review* 20:1–35.

Pater, Joe. 2006. The locus of exceptionality: Morpheme-specific phonology as constraint indexation. In *Umop: Papers in optimality iii*, ed. Leah Bateman and Adam Werle, 1–36. Amherst, MA: GLSA.

Pater, Joe. 2008a. Lexically conditioned variation in Harmonic Grammar. Presented at OCP-5, Toulouse-Le Mirail.

Pater, Joe. 2008b. Morpheme-specific phonology: Constraint indexation and inconsistency resolution. In *Phonological argumentation: Essays on evidence and motivation*, ed. Steve Parker, 1–33. Equinox Press.

Pater, Joe. 2008c. Serial Harmonic Grammar. Handout from Univeristy of Southern California Colloqium.

Pater, Joe. 2009. Weighted constraints in generative linguistics. *Cognitive Science* 33:999–1035.

Pater, Joe. to appear. Serial harmonic grammar and berber syllabification. In *Prosody matters: Essays in honor of Elisabeth o. selkirk*, ed. Toni Borowsky, Shigeto Kawahara, Takahito Shinya, and Mariko Sugahara. London: Equinox Press.

Pearce, Mary. 2003. Vowel harmony in kera. Master's thesis, University College London.

Pearce, Mary. 2008. Vowel harmony domains and vowel undershoot. In *Ucl working papers in linguistics*, volume 20. Research Department of Linguistics, University College London.

Pearce, Mary. 2011. Effects of atr, fronting and rounding harmony on reduction. Presented at the 19th Manchester Phonology Meeting.

Peuch, Gilbert. 1978. A cross-dialectal study of vowel harmony in maltese. In *CLS 14*, ed. D. Farkas, W. Jacobsen, and K. Todrys, 377–390.

Potts, Christopher, Joe Pater, Karen Jesney, Rajesh Bhatt, and Michael Becker. 2010. Harmonic Grammar with linear programming: From linear systems to linguistic typology. *Phonology* 27:77–117.

Prince, Alan. 2004. Anything goes. In *A new century of phonology and phonological theory*, ed. Takeru Honma, Masao Okazuke, Toshiyuiki Tabata, and Shin ichi Tanaka, 66–90. Kaitakusha.

Prince, Alan. 2007. The pursuit of theory. In *Cambridge handbook of phonology*, ed. Paul de Lacy. Cambridge University Press.

Pruitt, Kathryn. 2008. Iterative foot optimization and locality in stress systems. University of Massachusetts, Amherst. ROA-999.

Pulleyblank, Douglas. 1996. Neutral vowels in Optimality Theory: A comparison of Yoruba and Wolof. *Canadian Journal of Linguistics* 41:295–347.

Pulleyblank, Douglas. 2004. Harmony drivers: no disagreement allowed. In *Proceedings of the 28th Annual Meeting of the Berkeley Linguistics Society*.

Pulleyblank, Douglas, Ping Jiang-King, Myles Leitch, and Ọlanikẹ Ọla. 1995. Typological variation through constraint rankings: low vowels in tongue root harmony. In *Proceedings of the South Western Optimality Theory Workshop*, The University of Arizona Coyote Papers, 184–208.

R Development Core Team. 2009. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. URL `http://www.R-project.org`, ISBN 3-900051-07-0.

Rhodes, Russell. 2010. Vowel harmony as agreement by correspondence. Ms, University of California, Berkeley.

Riggle, Jason. 1999. Relational markedness in Bantu vowel height harmony. Master's thesis, University of California, Santa Cruz.

Ringen, Catherine. 1988. Vowel harmony: Theoretical implications. Doctoral Dissertation, Indiana University.

Ringen, Catherine, and Heinämäki. 1999. Variation in Finnish vowel harmony. *Natural Language and Linguistc Theory* 17:303–337.

Ringen, Catherine, and Robert M. Vago. 1995. A constraint based analysis of Hungarian vowel harmony. In *Approches to Hungarian: Levels and structures*, ed. István Kenesei, volume 5, 309–19. Szeged: Attila Josef University.

Rose, Sharon, and Rachel Walker. 2004. A typology of consonant agreement as correspondence. *Language* 80:475–531.

Schlindwein, Deborah. 1987. P-bearing units: A study of Kinande vowel harmony. In *NELS 17*, 551–567.

Smolensky, Paul. 1993. Harmony, markedness, and phonological activity. Handout to Presentation at the Rutgers Optimality Workshop 1.

Smolensky, Paul. 2006. *Optimality in phonology II: Harmonic completeness, local constraint conjunction, and feature domain markedness*, volume 2, chapter 14, 27–160. The MIT Press.

Smolensky, Paul, and Géraldine Legendre. 2006. *The harmonic mind*. MIT Press.

Sohn, Ho-min. 1975. *Woleaian Reference Grammar*. Honolulu: University of Hawaii Press.

Starwalt, Coleen. 2008. The accoustic correlates of ATR harmony in seven- and nine-vowel african languages: A phonetic inquiry into phonological structure. Doctoral Dissertation, University of Texas at Arlington.

Steriade, Donca. 1987. Redundant values. In *Papers from the 23rd Regional Meeting of the Chicago Linguistics Society, Part 2*, ed. Anna Bosch, Barbara Need, and Eric Schiller.

Steriade, Donca. 1995. Positional neutralization. In *North East Linguistic Society 24*. University of Massachusetts, Amherst.

Steriade, Donca. 2001. The phonology of perceptibility effects: The p-map and its consequences for constraint organization. In *The nature of the word: Essays in honor of Paul Kiparsky*. MIT Press.

Suomi, Kari. 1983. Palatal harmony: A perceptually motivated phenomenon? *Nordic Journal of Linguistics* 6:1–35.

Suomi, Kari, JM McQueen, and A. Cutler. 1997. Vowel harmony and speech segmentation in Finnish. *Journal of Memory and Language* 36:422–444.

Suomi, Kari, Juhani Toivanen, and Riikka Ylitalo. 2008. *Finnish sound structure: Phonetics, phonology, phonotactics, and prosody*. Studia Humanitoria Oulensia 9. Oulu: Oulu University Press.

Svantesson, Jan-Olaf, Anna Tsendina, Anastasia Mukhanova Karlsson, and Vivan Franzén. 2005. *The phonology of Mongolian*. Oxford University Press.

Terbeek, Dale. 1977. A cross-language multi-dimensional scaling study of vowel perception. Doctoral Dissertation, UCLA.

The Menominee Indian Tribe of Wisconsin. 2005. A beginner's dictionary of Menominee. Available online: http://ling.wisc.edu/ dictionary/testing/main/.

Vago, R. M. 1976. Theoretical implications of Hungarian vowel harmony. *Linguistic Inquiry* 7:243–263.

Vago, Robert. 1988. Vowel harmony in Finnish word games. In *Features, segmental structure and harmony processes, part ii*, ed. Harry van der Hulst and Norval Smith. Foris.

Välimaa-Blum, Riitta. 1999. A feature geometric description of Finnish vowel harmony covering both loans and native words. *Linuga* 108:247–268.

Vroomen, J, J Tuomaineen, and B de Gelder. 1998. The roles of word stress and vowel harmony in speech segmentation. *Journal of Memory and Language* 38:133–149.

Walker, Rachel. 1998. Nasalization, neutral segments, and opacity effects. Doctoral Dissertation, University of California, Santa Cruz.

Walker, Rachel. 2000. Long-distance consonant identity effects. In *Proceedings of the 19th West Coast Conference on Formal Linguistics*, ed. Roger Billerey and Brook Lillehaugen.

Walker, Rachel. 2004. Vowel feature licensing at a distance: Evidence from Northern Spanish language varieties. In *WCCFL 23 Proceedings*, ed. A. Kelleher B. Schmeiser, V. Chand and A. Rodriguez, 787–800.

Walker, Rachel. 2005. Weak triggers in vowel harmony. *Natural Language and Linguistc Theory* 23:917–989.

Walker, Rachel. 2009. Similarity-sensitive blocking and transparency in Menomini. Paper presented at the annual meeting of the Linguistics Society of America, San Francisco, CA, January 9, 2009.

Walker, Rachel. 2010a. Nonmyopic harmony and the nature of derivations. *Linguistic Inquiry* 41:169–179.

Walker, Rachel. 2010b. One-to-many relations. Paper presented at the 18th Manchester Phonology Meeting.

Watson, Janet C.E. 2002. *The phonology and morphology of Arabic*. Oxford University Press.

Wilson, Colin. 2000. Targeted constraints: An approach to contextual neutralization in Optimality Theory. Doctoral Dissertation, Johns Hopkins University.

Wilson, Colin. 2004. Analyzing unbounded spreading with constraints: marks, targets, and derivations. UCLA.

Wilson, Colin. 2006. Unbounded spreading is myopic. Bloomington, IN: PhonologyFest 2006.

Wilson, Colin. to appear. A targeted spreading imperative for nasal place assimilation. In *Proceedings of NELS 41*.

Wolf, Matthew. 2008. Optimal interleaving: Serial phonology-morphology interaction in a constraint-based model. Doctoral Dissertation, University of Massachusetts, Amherst.

Woock, Edith Bavin, and Michael Noonan. 1979. Vowel harmony in Lango. In *Papers from the 15th Annual Regional Meeting of the Chicago Linguistic Society*, ed. Paul R. Clyne, William F. Hanks, and Carol L. Hofbaur.

Yip, Moira. 1988. The obligatory contour principle and phonological rules: A loss of identity. *Linguistic Inquiry* 19:65–100.